Oral History of Fred Moore

Interviewed by:
Tom Gardner

Recorded: July 14, 2009
Boulder, Colorado

CHM Reference number: X5432.2009

[0:00:00] time stamp from beginning of DVD 1 of 2
Tom Gardner: This is Tom Gardner, a volunteer at the Computer History Museum here in gorgeous Boulder, Colorado interviewing today, Fred Moore, for historical products and things that happened at Storage Tech in the '70s, '80s, and '90s. Fred, why don't you tell us a little something about your background?

Fred Moore: Okay. Well my background begins in Crystal City, Missouri, south of St. Louis about 35 miles on the Mississippi river. I grew up there, went to the University of Missouri as a math major. I played basketball at the University of Missouri, was a math major and the head basketball coach there, Norm Stewart, said, "What's your major," one day. And I said, "Math." And he said "You're out of here. You've got no business playing basketball and being a math major, we don't have people like that here." So I went on after math at the University of Missouri. I did continue to play basketball but on some of the semipro teams around there. And got into math, and as a math student I got a job as a computer programmer at the University Medical Center and that paid my way through graduate school at the University of Missouri where I got very interested in computer applications and physical science and the earth sciences. That was in the early 1970s and about 9 years later that field of discipline became now as GIS, geographic information systems. I can remember doing my Master's thesis and I found eight articles that I could research that had ever done any computer applications in physical geography. At that time most geographers did maps and charts by hand. So I had a CalComp plotter at the University of Missouri when I got that job and I went wild to start to computerize all the things that people had done for centuries by hand. And it was a little bit of a rebellious Master's program I would say. But I had a great graduate school sponsor who let me do a hybrid masters in computer applications and physical science. When I went to interview for a job when I got my Master's degree, in addition to the degree you also needed to have work experience, which I didn't have in the GIS field, it was very new. And I wanted to come to Boulder, Colorado and work for NOAA [National Oceanic and Atmospheric Administration] or NCAR [National Center for Atmospheric Research] here. And neither would hire me because I didn't have any work experience. So then I went straight into the computer field and went to Dallas, Texas as Assistant Programmer at Republic National Bank. And we did all the internal operating system modification installations. I turned out to be the storage guy on that staff and I dealt with the storage vendors. And one day the Storage Tech vendor, who I'd gotten to know very well, came to me and said, "You know we're going to hire a new type of person at Storage Tech, brand new. We're moving from a tape company and becoming a disk company and we need somebody with your background who can explain how a disk works and help us sell it." Tape selling was more feature by feature a price, you know, match everything in the column IBM had. But disk was a lot different. And you know, need to do capacity planning, performance tuning, storage management, backup recovery, a whole different world. So I said, "Yeah, I'll talk to you about that." And Storage Tech in 1976 hired me in Dallas as the first systems engineer in the whole company. I did that for two years. I was the only SE in the entire company for two years. The company was $82 million in revenue at the time when I joined it. And I had all 53 salesmen in the world to support at that time.

Gardner: One system's engineer to 53 salespersons? Must have been a lot of air travel time?

Moore: A lot of air travel time, it was sort of a trial and error project. I mean, to be an SE was like an overhead. You know, we're going to pay this person to do something. Can we even measure what they do? And I traveled all over the world. I'd be gone about four days a week out of Dallas where I was based at the time and then the fifth day be in the area. And pretty soon a couple of years later they moved me to Colorado to build a worldwide SE organization.

Gardner: That's really amazing ratio. At other companies the SE ratio to salesmen was a lot lower. But going back a little bit in your background, you have any brothers, sisters?

Moore: I was an only child. My parents married late in life. My mom was a kindergarten teacher for 30 plus years and my dad worked at Pittsburg Plate Glass factory and he had a fourth grade education. So they were wondering about me you know. Could I, I thought I could hardly read, barely write, had a limited vocabulary and dressed like a flood victim but somehow I got through college and the university and graduate school and found my way to where I really wanted to live and that was Boulder, Colorado.

Gardner: Yes Boulder's just beautiful - where'd you go to high school?

Moore: Went to high school at Crystal City High School in Missouri, south of St. Louis. It was also the high school of our most famous citizen in Crystal City, a guy named Bill Bradley that played for the New York Knickerbockers and a Rhode Scholar NBA champion and all-star player. I knew Bill, played some ball with him growing up. He was several years older than me but still had some contact with him and he is on the board of Seagate Computers right now.

Gardner: So as SE in the mid to late '70s, that was just about the time that Storage Tech was getting into the disk drive business. What was that like?

Moore: Well the disk drives the reason for SEs. I did it for two years in the field and the company said, "Okay, this is what we need to do. Will you come to Colorado and build a worldwide force?" So for the next five and a half years we spend building up to about 240 SEs around the world to support worldwide sales of discs. And a lot of innovations that was to come in the Storage Tech product lines, so a lot of time finding people, bringing them in, training them, making them storage experts on disk tape and the entire storage hierarchy. And a lot of people back then always looked at Storage Tech as tape technology. And what I liked about the name of the company was storage technology. We were disk, we were tape, we were software to help manage the storage and the word storage was the one that intrigued me. I was always interested in all aspects of the storage hierarchy not just one specific product line.

Gardner: That era was the era of the 3350 and the double density, what was it like introducing those products into the customer base?

Moore: Well the 3350 family of products of course was IBM's first entree really into fixed media disk. And Storage Tech had already been there with a prior product call the Super Disk, it was non-removable. So we had a little bit of a lever to go tell people about how you use fixed media disk. People liked to remove discs. We'd already learned how to remove data that was non-removable and like that. And we needed to move that Storage Tech off of the 3330 type format that the Super Disk was on because of a large number of field reliability issues. So the new family of products called the 8350, comparable to IBMs 3350, was a very exciting time for us. I can remember as an SE, our whole game was to find an advantage on our product line over someone else's, not just match by price and feature but what is your real value proposition? And that was the secret sauce for SEs and marketing people, was to make your product different and more valuable than the competitors in some way, some set of features, some set of function that could separate you from the competition. Those were very, very exciting times.

Gardner: What was your secret sauce for the 8350?

Moore: Well the 8350 product line initially came out to be compatible and it had a price advantage but soon after that the innovation at Storage Tech really did begin to kick in. We introduced a double density version of the 3350 that IBM did not have. Storage Tech pioneered a concept on disk called Duel Port which allowed two simultaneous data transfers to a storage subsystem where IBM was still using at the time in architecture called String Switching, if anybody can remember those terms. So we had a higher through point number and performance became the differentiator on that. And then secondly, when you had double density you had twice the amount of data in the same physical footprint. You had a storage density improvement on gigabytes per square foot which played into data center, costs, occupancy, and those things. So the story began to build and those were the early seeds of innovation at Storage Tech to make things compatible yet different and more valuable than their competitors.

Gardner: There's no question that the 8350, 8650 were huge successes. There was a problem in '82 when apparently there was a huge recall of that product family? You recall anything about that?

Moore: <Laughs> I do recall the recall and the recall was primarily on the 8650, the duel density products. We had built our own media for those products. The media was basically not that reliable. In fact reliability was a differentiator at that time in favor of IBM. And the 8650 reliability became a major challenge for our field engineering force to keep up and the company struggled for two years plus to get a very hard and reliable media in the product. New versions came out all the time and put a lot of work on the field sales force and the field engineering team to work with the customer, try this, try that. Frankly the company was in trial and error mode to try to get the media stabilized for the product. Some customers had no trouble but by in large the percentage of failures in the install base was great enough that the product actually cost the company money with the rework.

Gardner: I didn't realize Storage Technology made its own media? Is that something that was introduced with the 8650?

Moore: Yeah, the company did make its own media for a while and obviously began to realize that other people were better at that than Storage Tech was, so they went outside and began to bring in media from other companies like Dysan, for example, is one name that stands out that the media came in and finally began to get those products cleaned up and get positioned for what was the next wave of products which was the 8380 type family compatible to IBM's 3380. But those were a rough few years and in my view though the disk issues and the cost of the work on the disk reliability were the single primary cost for the company going into chapter. There were other causes for sure but at that time that was the biggest one and it put a lot of burden on the company to work with the customers to keep them believing that we could ultimately produce a successful program.

Gardner: Do you have any idea what the underlying media problem or problems were?

Moore: Well it dealt with the hardness of the media. And you know the media, they could just never get the, as the engineers would say to us, they could never get the formula right. It was like baking a cake at different altitudes. One day would be good, one day it wouldn't be so good. Your barometric pressure

changes, things were a little off, it was that tight in there and the process was as much as an art as a science. So bottom line is they never mastered the process.

Gardner: I think most of your OEM suppliers and the media at that time were at sea level and we're not at sea level here in Louisville.

Moore: Yeah, we're a mile high here and people even brought that up, that all those factors had to do with that. And...

Gardner: Was the media first introduced with the 8650 or was it earlier products?

Moore: 8350.

Gardner: And then...

Moore: And then the 8650 put more pressure on that because the tracks were much closer together so therefore the coercivity, and the you know you had signal destabilization because the tracks were close and magnetic fields interfered with head crashes. And as an SE manager, I wound up spending a lot of my own time out working with customers to help them through that. And by the way, this set the stage, this level of reliability problem, set the stage in the industry for a concept that would appear about six, seven years later called RAID. Disk reliability was not great even at IBM. It was better than what Storage Tech had but the RAID concept was designed later on to allow for a disk substance to stay available if the disk crashed. In the 8650 world if a disk crashed, you're done. Your data's gone, your applications die, maybe your whole system. But with RAID systems later on, a disk could crash and it had enough carry to recover from that failure to keep your business alive.
[0:15:00]
And that was a huge plus that by the close of the 1980 decade was one of the most positive steps the disk industry ever took.

Gardner: This time we're now still in the late '70s, early '80s. Storage Tech began diversification, a plethora of new products, new innovative products, I'd like to move from the disk storage or from the IBM compatible, directly compatible storage to some of the others. For example, what can you tell us about the 4305?

Moore: Well the 4305 product was really very innovative and a few people at Storage Tech headquarters, a guy named Andy Anderson, who still lives here in Boulder and a guy named Gary Holtwick kind of drove the marketing effort to go and see would a customer pay a lot more per gigabyte of storage for a device that delivered 1000 times the IO performances disk with consistent response time every time? And by the way, the product was non-volatile; it was volatile so if the power went away, it forgot everything it had in it. And that was, they spend six months on the market research to see if a product that had those capabilities would sell. And people said we need performance for certain data sets in our business. This was a mainframe product and certain data sets at the time were the system catalog and job queue and things and paging for virtual storage paging that were extremely IO intense

both for reads and writes. And the market research came back and I get to survey a lot of customers as part of that myself because I was in Dallas at the time that happened. And people said, "You know what? It's good." And that product came out in 1978 at $8800 per megabyte list price. If you, today we price all storage by gigabyte not megabyte so that's $8.8 million a gigabyte. So just to compare that to SATA disk arrays today that are $1 to $5 gigabyte.

Gardner: You mentioned two names, Andy Anderson
and...

Moore: Gary Holtwick, H-O-L-T-W-I-C-K. He was the Director of Marketing at the time.

Gardner: Do you know any of the technical folks who did the product architecture?

Moore: I do, I'd have to dig and think about who their names were on the solid state disk.

Moore: One guy, his name was Erik Rinkjob, and Erik was an engineer doing a lot of work on the solid state disk at the time.

Gardner: SSD was sort of a leading edge product, right?

Moore: Well it was completely innovative and it was also a virtual storage product. It looked like disk to the operating system. In fact, in this case it looked like IBMs 2305 fixed head disk, also known as a drum. And so it virtualized that, it was really DRAM chips, not a drum but it made the DRAM chips look like a drum. So you have really Storage Tech was a major pioneer in this time in virtualizing storage. So this was a great example. You opened up the box and saw DRAM chips, the system saw a 2305. In fact it saw multiple 2305s because the capacity of that product went up from 48 to 192 gigabytes.

Gardner: As announced or over its life.

Moore: Over its life.

Gardner: It was a long life product too.

Moore: It lasted; the 4305 was announced in '78 and delivered, not any moving parts. It was really, people loved that no moving parts thing and it really lasted until 1985. And on IBMs 3090 mainframe family they introduced a new feature, which was also DRAM chips. It was an extension of main memory called expanded storage. And expanded storage allowed all those paging and swapping files, that the mainframe used, that was the primary application for solid state disk.
[0:20:00]
It allowed all of those to reside in expanded storage. So now I had an extension off the memory bus to contain that data as opposed to a device out on an IO channel, like a solid state disk. And the expanded

storage was even much faster than the SSD and it was also cheaper. And that was really the end of the SSD world. And the reason the SSD couldn't survive, because it was volatile. With paging gone now you're putting permanent data sets on a volatile product, so when the power went away the data went away. They did have battery backup and other things, but people were still very concerned about that so the product could not survive on its own, though it would reemerge and I have a reemergence here twenty some years later of solid state products in a brand new form, by the way. But the product barely kept alive after 1985. Two to three companies, small companies for non-mainframe systems had small ones out there and it never did come back on the mainframe.

Gardner: So basically as the new processors displaced the older processors the machines came off rent or were scrapped.

Moore: Scrapped.

Gardner: Do you think the same things going to happen today with SSD?

Moore: I don't think the same thing will happen today with the SSD and the main reason is that today's SSDs don't use DRAM chips, they use flash memory. So they are not volatile, you can pull the plug and the data stays intact. They're not nearly as fast as DRAM but they are 25 to 50 times faster than magnetic disk, as opposed to 1000 times like DRAM. So these devices are, right now as we speak, entering all phases of computing. Almost every storage vendor has some type of flash product offering. They emulate a rotating disk in a two and a half inch disk form factor with flash chips, plug it in, plug it out, you open it up and there's flash chips in there and it's much faster. And that is the new wave of non-volatile SSD. They are nonvolatile. They are removable, which magnetic discs are you know it is very tough to remove magnetic disk. You can do it, but it's a risky venture. So the new SSDs I think will stick. They're slower, they're dramatically cheaper. In fact the price of the flash memory SSDs are coming down today right on top of cache/disk subsystems.

Gardner: Actually I was struck by the mainframe analogy where the, it wound up being better to put the extra memory up on the memory bus instead of IO bus. And there are people in the SSD space today who are saying, "No flash will wind up on the memory bus as memory more of it either directly addressed because of the addressing is large enough today to handle it or virtualization is in some way." And, but you know that's where you want to see it, not out on the slow IO bus. The analogy would be IO Fusion, the company that is pitching their product that plugs into a PCI bus, a memory bus product rather than an IO product. But if I could predict the future I'd be a lot wealthier than I am.

Moore: Yeah, well maybe you can on this one. Because if you put it on the memory bus you put data in that SSD that's only unique to that server or that operating system, but to share that with other servers is very difficult. If you want to share the general population of files and data sets you put it out on the data channel. So for system use it's okay to do that on the memory bus. But you would never put your database on the memory bus, you'd put your database on a channel in case one server went down another one could get to it. So it depends on the application being used. System files bring it as close to the server as possible but for general purpose applications, which is the vast majority of disk like data, you want that on your IO channel. So two different, completely different used for that product.

Gardner: Any interesting application stories of installing SSDs, any stories interesting both in successes and failures?

Moore: Yeah, well I guess the greatest thrill for an SE ever was probably to put in a SSD, take the measurements of paging and swapping response time before and then take it after and watch that customer go. "Wow." I mean this was one thing where if you could put a price on the performance improvement you know, what is this worth to your business to be 50 times faster or 1000 times faster? It would be great. But I think those are the great moments and they happened a lot. The bad moments were the power went out in the building. We had a power outage or a hurricane hit, a flood, whatever it was. The power went away and I lost all my paging files and my system crashed and now I've got a 24 hour recovery ahead of me. And that prompted Storage Tech to put in battery backups that allowed power to stay on the DRAM chips long enough to drain all the data off the DRAM into fixed disk in to the device.

Gardner: Oh really?

Moore: So that became the solution for that. The sad part of that was that when you put all that extra hardware in there, the price of the SSD went up even higher because now I had disk in there and I had Microcode and I had Logic and I had everything to drain it. So the price point were way beyond disk and most customers, the hardest thing with SSD over all was to convince a customer that the price, that much more than disk, were worth it. They knew it would be faster but is that much performance worth that much more price.

Gardner: Now you gave us the absolute price at the beginning at this section, what was the relative price then? So how much, do you remember what an SSD, a 4305 was versus say a 8650, in terms of dollars per megabyte not gigabyte?

Moore: Well disk drives you know would go for over $1000, $1200 a gigabyte in that time frame. So it was 7, 8 times more than that.

Gardner: Yeah that's more or less where SSDs are today, I mean 10 times.

Moore: Yeah they are.

Gardner: Comparable.

Moore: But the absolute dollars were staggering for some people to deal with back then. And the price of discs then had begun to fall, you know, at 30 points a year.
[0:26:51] Start of = VTS_01_2.VOB
And I think one of the things about pricing, list prices in the storage industry are like a stop sign in Italy, they're merely a suggestion. So you would see discounting you know 40, 50 percent off on them. The solid state disk had tremendous margins. It's the highest margin storage product I've ever been

associated with. Those margins you know would hit 70 percent at times on that. So there's great margin there and the sales force had a lot to work with. But these were the days when price became the final answer and people put value in second place. And I was always concerned about that as an SE this might cost more but its better and worth more. And people couldn't quite make that step, they had sticker price fever, you know it's like you know buying a new car. You go out and buy a new car, you get a great deal. At the end of the year, you come home and say I got this great deal but you know three months later you've got insurance, taxes, emissions fees, license fees, maintenance fees, you're buried in all the other stuff. And this industry of course is struggling but finally and slowly beginning to move away from sticker price and looking at the overall operating cost of something.

Gardner: About that time Storage Tech came out with a drive not in the mainframe space, the 2700. Did you in your SE role have much to do with the 2700?

Moore: You know, we did not have a lot to do with the 2700 those were generally products going through a different sales channel than the direct sales force. They went through OEM channels. They would go through nontraditional sales products. So they were products engineered and a little bit out of the SE realm. The SEs were direct sales support.

Gardner: Back on that 4305 the, it had a model number change before. There were several different models that came out until the last model which was about 1985?

Moore: You know that's a great point. With the 4305 was named after IBMs 2305 drum so it looked like a drum.

Gardner: The operating system thought it was one or more 2305s.

Moore: 2305 Model II drums. Now the 4080, so we're a company we can innovate in storage, we're getting really good at virtualization, we said, "Wow, IBM's announced the 3380, let's make this look like a 4080 so that it doesn't look like the drum anymore," because IBM dropped support out of their old drum. So the system didn't do it so now with the 4080 we made the, what was the 4305 look like a 3380. We called it the 4080 so know it looked like a disk, not a drum. And it looked like a very, very fast 3380. You know the operating systems, here's a 3380 and there's one, there's one but man, this one is a 1000 times faster than all the rest.

Gardner: Did it, I think the 3380 minimum capacity that was addressable system was either 600 megabytes or 1.2 gigabytes.

Moore: No, 630 megabytes per actuator, the 3380. So the 4380 was smaller per actuator. These were smaller virtualized versions.

Gardner: So you somehow had to deal with this was a 3380 volume but it wasn't full capacity.

Moore: Yeah, whenever you jammed up your system, the 4080 looked like a 3380 but what we would do is present the fact that 75 percent of the capacity was already allocated so all you had was a fourth of it left to use.

Gardner: So that was part of the installation with you folks.

Moore: Microcode did all that for me.

Gardner: Wasn't, it wasn't required that it be allocated at the system level.

Moore: Microcode set it up for you so that you had 100 plus megabytes of storage that you could use right there.

Gardner: And I presume it was, had varying numbers of storage. Small, large, had a range of capacity?

Moore: You could set them up on different sizes and vary them around and configure if you wanted, you bet.

Gardner: And that was the 4080?

Gardner: And that was the last of them.

Moore: That was the end of it. And that one, we'd already been hit with expanded storage and now to position the 4080 with the battery backup and the disk reserve to drain that data off DRAM in there and to sell it as a high performance disk was the last stand really for the solid state disk in this era. And the problem there was then it became a price issue. You know, paying so much more for this disk than that disk. And you were fighting the fact that rotating disk was falling at 30 percent or more per year on a per gigabyte basis.

Gardner: It was also a cache product in that time period. How do you differentiate the cache product dropped from the SSD?

Moore: Well with a solid state disk, data's permanently allocated in the SSD and it stay there. With the cache, data stays in the cache only for a relatively short duration of time. So the cache was really a buffer for backend disk where the solid state disk was its own entire storage device. Data stayed in there for all activity. So the cache product came out in 1981, it was called the 8890 Cyber Cache. The disk program manager for that product actually works in this office with me and he's my mountain climbing partner, Randy Watson who you met earlier, and the cyber cache came out in anywhere from 1.5 to 18 megabytes of cache to frontend the backend of the subsystem. The frontend of the 3350 compatible line for Storage Tech and IBMs first cache came out and they only offered cache on their 3380. So our differentiator was, we offered cache on the 3350 class products. But IBM didn't bring cache out until the

3380 came out. So that was a real plus for Storage Tech because now we could offer high performance data on a cache. If you read data off the disk the data gets read into the cache and all subsequent accesses come out of the cache for that data as opposed to going back to the disk all the time. And therefore, you know, repeated use, which we used to call locality of reference, of all the data this much of the data was the hot data. That's what stayed in the cache, and that got all the access. So it perked up a disk subsystem considerably. And we would typically find 70 percent to 80 percent of all the IOs for read access, read to a disk subsystem to come out of cache. That's at DRAM speed versus rotating disk speed so, big plus on that one.

Gardner: Now was that a product that sat in the channel, between the channel and the control unit, or was it a modification to the control unit for the disk drives?

Moore: The 8890 cyber cache was a modification of the control unit.
[]
It was a control unit feature. People would order control units with or without cache at that point in time. Cache was not standard like it is today in every disk control unit. So people would do that, the systems engineers would configure how much cache they needed as if that was going from 1 to 18 megabytes was a big deal back then. Today that would be nothing but they went, "Oh I need 12, I need 18, I need 6," you know, and it was really finally tuned that way.

Gardner: So you took the control unit for the 8830 family, the 3380 equivalent family, the cached control unit is not the control unit for the 8350 and 8650, it's based upon the next generation control unit, the 8880.

Moore: For the 8380 disk family. And that became a feature in there and the cache got, then cache got much bigger. The backend disk where much bigger, there was a lot more data to cache. You know people would read data into a lot of accesses off of data. It's not like I sequentially read everything once and I'm done. But data would come in and get hammered for two or three hours on databases and this is kind of when this storage really took off in large capacity subsystems with the 3380 family of products, which by the ways a longest lasting family of disk in the history of the disk industry.

Gardner: You said it supported the 8650, 8350.

Moore: The 8890 was the first one.

Gardner: The 8890 so the new, the 8890 control unit or cache control unit, supported the older disk drives and also the new 3380 compatible, the 8380, did I get that right?

Moore: The 8890 supported the 3350 compatible class of disk. So Storage Tech had cache before IBM.

Gardner: Okay.

Moore: IBM introduced cache on their 3880 controller, two of them, a mod 11 cache for paging, believe it or not, and a mod 13, which is the general purpose cache. So in their 3880 controller, they had cache. When we introduced the 8380 compatible disk to their 3380 disk we put cache in that control unit. So the 8890 was for the 3350 family.

Gardner: That's what I was trying to get out.

Moore: There's a new control unit for 3380 all together. It was duel ported. You know it had a whole different game.

Gardner: So the 8890 control unit was a redesigned control unit for the 8350 family basically allowing a large cache to be installed.

Moore: A midlife kicker for that control unit was the advent of cache buffering.

Gardner: Is that how it got generated as a product? It was intended literally as a midlife kicker, is that the genesis of the idea?

Moore: Well when the control unit first came out, we didn't realize that somewhere in its life we would want to add a cache. But as word got around, IBM was going to work on a cache, a new family of product you know we knew at Storage Tech 18 to 24 months in advance that they had a new disk geometry coming. You know 47,000 byte tracks as opposed to 19,000 byte tracks, 3380 family, we knew that. So we knew cache was coming and therefore we added that to the current product line.

END OF TAPE 1 / BEGINNING OF TAPE 2

[0:32:59] Beginning of tape 2, in …3.VOB
Gardner: So the 8890 cyber cache, it probably lasted as long as the 8650, 8350 product line and as those were displaced by 3380 class the product went off ramp. Was it a successful midlife kicker?

Moore: Oh it was definitely. That was a good kicker and the company needed that kicker, you know, because performance became a better story than the reliability issues that they'd been going through. So it turned out really well. People liked cache, they loved the performance and from then on, cache became a standard feature in almost every control unit shipped. The option was not to have it but, you know, here we are 30 years later and nearly every control unit in the world shipped is cached.

Gardner: You just don't build it any other way.

Moore: You don't build it any other way.

Gardner: Now, were there choices of algorithms in terms of writing? Was it a true write through, write back type of control unit? What were the capabilities from a caching point of view?

Moore: So the cache cached read operations.

Gardner: Only read.

Moore: Only reads. So it-- if you wanted to read a sequential file you had some options and you could turn that on by file-- you could do sequential read ahead which read the entire file in, or you could read three or four blocks in advance and to try to stay ahead of what you were doing sequentially. So you could read the whole thing in or a few blocks like that. With random access, data would stay in, you know, record one, record nine, record thirty, record fourteen, they would stay in the cache until the LRU, the least recently used themselves out of the cache. They weren't used enough they hadn't been touched, something else needed to come in and they would get demoted. When data got demoted back to disk, if it had been, you know, the data-- all rights let me put it this way; anything written to the disk subsystem would write right through the cache to disk. There was no capability yet in disk control units to do what we call cache fast write, where there was a memory capability, a nonvolatile memory in the cache to cache the write operations. Later caches allowed writes to occur to the cache control unit and they would go to DRAM memory, nonvolatile and that would post to the operating system that the data was successfully on disk when it really wasn't. And then later on as soon as possible basically you would move out and that updated record would get written to disk. But the early caches were read only caches. And that was okay because 70, 75 percent of your IO activity was read.

Gardner: So when did the Storage Tech implement a write behind algorithm?

Moore: That was with the 8380 caches. The-- which was compatible with the IBMs 3880 model 13 cache, when that cache came out they offered two versions. One for paging called the 3880 mod 11, which was kind of a cache for paging that never went very far. But the big one was the 3880 mod 13 for general purpose data sets and files. And that one had cache, fast write, and then for all 8380s at Storage Tech, cache fast write was also implemented there. But the 8890, the prior one for the 3350, 8350 product line, those were read only caches.

Gardner: So Storage Tech then followed IBM's 80, 30, 3880-13 with write behind caching. What was the term you used, write?

Moore: Cache fast write.

Gardner: Cache fast write.

Moore: Right. So it was cache and the write was fast because it happened in DRAM and the system thought it was on the disk, it really wasn't. So you had to then trust that system that it would get out there to disk pretty quick and that worked okay.

Gardner: Well, those must have been in the mainframe world, know there's a whole different language between the mainframe world and the OEM world, like cache fast write, being the way IBM probably announced it in their announcement letter and then describe it in their feature material and it's used in the mainframe. I think in the rest of the world we would say write behind.

Moore: Write behind <laughs>, yeah.

Gardner: Right, yeah, write behind I think is the term that I'm more familiar with but...

Moore: Yeah that was their announcement terminology.

Gardner: DASD is another example.

Moore: DASD, you never hear the word DASD anywhere anymore, even in the mainframe world.

Gardner: Even in the mainframe world it's all disk.

Moore: Right, its all disk.

Gardner: Even if you're running whatever the latest operating system that's-- it's disk and not DASD.

Moore: If you said DASD today the audience would look around and there might be some old guy in the back of the room going, "Yeah I remember that."

Gardner: So we're dating ourselves.

Moore: We are dating ourselves here.

Gardner: Okay. Do you know who was the architect, who, from an engineering or even a marketing point of view, who were the key players on the cyber cache?

Moore: Well Randy here.
[0:50:00]
Randy Watson was the marketing manager for that product. He was the disk product manager for the company and led the charge on the 8890 marketing team. Another guy that followed him, his name was Jim Fleming, he was the disk product manager that followed after that and then a guy named Tom Purtell, P-U-R-T-E-L-L; became director of all disk product marketing disks, solid state disks and everything else. So that was kind of the sequence of people in the middle 1980 timeframe that led the disk charge at Storage Tech from a marketing perspective.

Gardner: Can you think of any of the key engineering guys?

Moore: A guy named Bob Hancock was a huge player in those products. He was an engineer, very performance savvy, he-- you know, one of the things about engineers they love to say how much the customer does this or that and they never talk to customers. So you always have to be very careful, but Bob was one that did talk to customers and we used him with customers. He was very presentable and understanding and it really helped to get customer input in that. So Bob Hancock was definitely one of them. They key engineer by the way on the 8890, his name comes to me now; his name was Bill Chenworth.

Gardner: C-H-E-N-W-O-R-T-H?

Moore: Right. He was the man. I'm digging it out of my memory bank here, but he was clearly the guy who led that whole charge like that.

Gardner: Any recollections on the 4305 class, now that you've had one set of brain cells?

Moore: You know, I'm working on that as we sit here so hopefully I'll dredge that name up, I can't believe I don't right now, but Hancock may have had a big hand in that one also at that time.

Gardner: Okay, another interesting attempt in this time period, now into the early '80s is something called VSS or virtual storage system?

Moore: VSS was a extremely innovative concept at Shortage Tech and I mean, I have to say personally, one of the great things working at Storage Tech was we were taking IBM's base architecture in building on it, enhancing on it, creating new ways to do it, you know, to make rotating disks look like high performance disks, whatever. The innovation was extremely good at Storage Tech, and VSS was the first of two products to epitomize innovation using IBM as a base architecture yet going way beyond that.

And the VSS as it suggests, virtual storage subsystem, virtualized storage, which Storage Tech was doing in a number of products and it integrated together a hierarchy of technologies. You've seen the pyramid in the storage industry-- in fact, I gave you one here earlier which is that hierarchy and it moves from high performance at the top of the pyramid to high capacity at the bottom. And what the VSS was going to do was take that pyramid and integrate it together, glue it together with software that would promote highly active data to the top of the hierarchy and demote less active data down and manage the location of that data based on its performance, based on its activity, its recall, its usage or whatever and find the optimal place to store it so you could optimize your storage infrastructure. Without a concept like that, people had to put everything on disk and half of what was on disk is a generalization, probably should have been on a tape, because it wasn't used enough. You know data that's not being used doesn't need to sit up there and consume power 24 hours a day if it's not being used.

So the VSS was going to build on a concept that IBM had introduced in the mid '70s called HSM; hierarchical storage manager. It was a software product and it moved data from disk off to lower cost storage, namely tape, based on activity, usage, LRU least recently used algorithms, or user defined policies that you set up. And the VSS's goal was to implement HSM like software into this storage subsystem to integrate solid state disk, cache disk, rotating disk, and tape, four layers into a hierarchy. And this was going to be-- you put your data there and basically you said that's it, I don't know where it's going to be, but it'll be where it needs to be because it's getting hit a lot, I'm going to be-- it's going to be up in high performance where it's not being used much, it'll be in low cost storage. And the whole goal was to optimize that.

And the company shipped a few beta sites on that product, very complex for that time. In fact, even today, in the year 2009 we sit and we look at an industry still trying to do something like this. This is almost 30 years later that the industry has not been able to implement what this does. Some people are trying it; there are a few versions out of things that move data around a little bit like that. But if that product came out today it would be an advanced architecture 30 years later. And basically, the bottom line was, it was too complex to make work effectively. There were some sites, a handful of beta sites, namely in California as I recall, and they just could not get everything to work well together. The product was eventually scrapped for financial reasons. It was too tough. Here we were asking a team of, you know, eight minute milers to run a six minute mile. And this was a great challenge and it was too much to do.

Gardner: My usual question, key people involved both from a marketing concept point of view and an architectural point of view, could you recall any names for us?

Moore: Well I can on this one. And the lead guy on this program's name was Barry White. Barry White was the man, this was his vision, his dream, his architecture. He had some help from some other guys like Jack Scott, a fellow named Mogens Pederson.

Gardner: Now you're going to have to spell that for me.

Moore: M-O-G-E-N-S P-E-D-E-R-S-O-N, a guy named Bob Raicer, R-A-I-C-E-R, and the notion here was sensational. The appeal was big, it was not over marketed or over hyped, which would come later in the history of Storage Tech. But the product was-- had tremendous appeal and people said I need this. You know, great play was the movie industry where the movies would be used a lot the first three or four weeks and then they'd go into basically archival status. Why not move that data down automatically where it would be safe, which was tape. You don't need to leave it on disk where it's not being used. So there were some great applications, everything made sense but the complexity here was like getting that first man on the moon, you know, it took a long time to do that.

Gardner: Since apparently the system was to be unaware that the data was being staged, it must have appeared to the system as a large pool of virtual disks?

Moore: Large pool of virtual disks, exactly. All data was cataloged and it might appear to be on disk when it was really on tape, exactly, exactly.

Gardner: So if the system thought the data set was on catalog to this DASD volume and it accessed it across an IBM channel to a storage control unit, but where it physically was was somehow controlled in this subsystem.

Moore: That's exactly right. And that's one of the fundamental premises of the HSM concept. Even today, HSM on the IBM mainframe works that way.

Gardner: The system software now moves the volumes physically from the disk to the tape, it's not hidden.

Moore: Right, this was going to all happen outboard in the storage architecture away from the operating system.

Gardner: So the catalog changes today in an HSM system. A data set is cataloged with tape or to a library or even to an archive or it's on a sub-- on a disk storage subsystem with its caching. It's fairly limited and it is at the system level as opposed to the invisibility of VSS. So could you tell us a-- was this then a subsystem of conventional storage type products, the tape drives were Storage Tech tape drives and the control units were disk control units in the cache. I mean, how was this then architected physically and logically to happen? There was a master control unit that sat above conventional tapes and disks and moved the data around or on a record basis? I mean I'd like to understand more about the architecture if you can.

Moore: So the box was more than just a box that linked together your hardware on the floor. It was a new hardware architecture.
[0:41:03] beginning of VTS_01_3.VOB
So within that device you had DRAM components that borrowed a lot from the solid state disk product. You had a cache disk subsystem which was essentially the guts of the Storage Tech disk subsystem. Then you had Storage Tech tape. But it was packaged differently, so it looked like a brand new system, it wasn't cobbled together with existing product lines. It looked like something brand new.

Gardner: And so it had in it a DASD control unit.

Moore: Yeah, but it had borrowed from a lot of the things that were in place, obviously. It would be a financial disaster to recreate everything and Storage Tech had all the building blocks there except the software.

Gardner: And the internal transfer was...

Moore: This was a software challenge. It had many internal data paths to move data back and forth, it were all architected around that so you were not going to see a lot of staging delays or demotion delays to a lower level on the hierarchy. But it borrowed and used a lot of the existing technology that was on

hand at Storage Tech. You know, remember at this time Storage Tech had a rich solid state disk product line, a rich disk product line, tape, I mean they had a lot of things to use.

Gardner: It's pretty interesting actually.

Moore: And by the way at that time the tape piece was still a nine track tape. We hadn't had the announcement of cartridge tape yet.

Gardner: So, and since you used human beings to mount.

Moore: Still had to do that.

Gardner: So it was essentially unlimited capacity, as many tapes as you could stack in a library.

Moore: Virtual, you never ran out of space.

Gardner: And how'd you then cause the mount command to get to the console when, you know, the system thinks it's on DASD and it's not and you know you've sent it to a tape that's in physically not in your control, how is that handled in the architecture if you know?

Moore: When the data left disk and went to tape it left a stub or a pointer. And that stub, you know, whenever you went back and said, "Whoops, it's really on tape." Then you would have an elongated disk access to get the tape mounted again like that. So you had to leave that stub there.

Gardner: Oh no I understand that, but I'm just curious as how you got the console message to print out at the system to tell the operator to, or better yet the message to print out of the library to tell the operator, to tell a operator to mount volume such and such on unit such and such, that's what's got to happen, right? You know about it in the VSS but MVS or whatever operating system's running doesn't know about it.

Moore: One of the great things about the mainframe operating system is that applications could issue commands. So it could issue its own mount commands, it could issue its own WTO; write to operator, WTOR; write to operator with reply command so we could issue all kind of commands and did to the operator console to generate mounts and demounts and all that stuff. So that was part of that software.

Gardner: Okay, so there was an application running at the host level, which communicated to the VSS and thereby those types of commands.

Moore: **Y**es, absolutely it generated all those operator commands. You had to do that.

Gardner: MVS or whatever the IBM operating system was had no idea that there was this mount taking place, it was a communication from VSS to the application, to the physical mount.

Moore: And really that's a concept that was used later on when the Nearline tape library came out, you had the same thing that you had to do, you know, with a mount command and the robot had to go get it you had to have a host software component running that managed mounts and all that sort of thing.

Gardner: Well in the IBM operating system managed mounts when it knows the volume is not on a tape. It will do that.

Moore: Right.

Gardner: And you can redirect that to the library which I guess what you do with a library system so that there's no physical operator and the mount goes out to the library and the library mounts it. Same idea though.

Moore: You needed to add to the basic IBM architecture with your mounts and dismounts and where your data really was in the hierarchy to go find it. But you still wanted to make the IBM operating system think it was just in one place. But what you did in the background to go find it and retrieve it, that was up to you.

Gardner: But the underlying drives were the drives that the operating system viewed as virtual drives. So if you had a pool of physical 8380s out there.

Moore: Yes

Gardner: You weren't-- the physical records were not being virtualized.

Moore: Well I'm not sure if I understand that or not. But you had this much physical disk, but because you had so many virtual versions it looked like it was this much disk. But just some of that would be paged out off on the tape or something like that, so. But it all appeared to be on disk to the system. Some disks just took a lot longer to get at those records than others.

Gardner: Now what I'm addressing is the-- you have a data set that's organized in a-- under a certain VSAM, ISAM, a certain access method which divides it up into a certain number of records with their counts and their keys and their data. And that's records that are developed according to the access method of the device; 47,000 byte tracks on an 8380, 19,000 byte tracks on an 8650. Whatever it is that's all handed. But if you're now appearing as a virtual 8680 you're still a 43,000 byte track and that track is broken up into whatever records with whatever count key and data lengths you've set up for the operating system and you're not breaking that record up and making it look like a bunch of 512 byte records. The next step in the evolution when we get the rate...

Moore: That came later.

Gardner: I'm trying to understand the...

Moore: Yeah, we're not at Iceberg yet.

Gardner: We're not at Iceberg yet, that's what I was trying to understand. And why we're...

Moore: Iceberg did what you described. It would change the data into chunks. Fixed block chunks actually.

Gardner: Yeah, the next, you know, those of us who've been in the mainframe space know and love count key data, but most of the world hates it, or denigrates it I guess is the best way to say it.

Moore: That's the best way to say it.

Gardner: It certainly has had its experiences. Anything else on VSS that-- you said there were a few-- you-- do you recall the sites, actual beta sites?

Moore: One was Hughes Aircraft. And I'm thinking of another in the movie industry out there but I'm not sure right now. There were under five shipped. And they could-- the program couldn't get out of technically beta testing. And it was ultimately killed because it became a cash drain. Remember the company was beginning to fight disk reliability problems, you know, and some of the bottom line expenses were going back in and the company let go of some of those things.

Gardner: We'll talk about the bankruptcy and the turnaround. What were some of the beta problems that stand out, if any about the VSS -- just too many lines of microcode or?

Moore: Too much code, you know, the code didn't work as advertised. There were just a lot of code to integrate and it became a major integration challenge to get all those pieces to work. And here you're asking a storage subsystem to recreate most of the HSM function that ran very cleanly on a mainframe. That's a lot of stuff to build into combination of microcode and a piece of host software. And it was the complexity. I mean given enough time, this thing would have come out. And a number of people in the company argued mightily to give it more time give it another year or so to come out. And the debate within the company at the time was big. I mean the, you know, it was an aggressive debate and a hard fought one to finally stop that program. It didn't go down easily. Everybody didn't celebrate when it was cancelled. And it had well over 100 people working on it at one point in time.

Gardner: Had you moved up to VP of Marketing by then or you were still in systems engineering?

Moore: It came after Chapter 11.

Gardner: Okay, so you're...

Moore: I was still an SE Manager then.

Gardner: Still the SE Manager.

Moore: And we were, as SEs at the time waiting for this, this was going to be the coup de grace for an SE. Now we had everything integrated into one great subsystem. I mean SEs loved things that were not PCM compatible because that was a feature by feature price, boom here's my-- but these other things were unique with different advantages and we were all pulling for the program to be successful. But I'd certainly realized that you couldn't fund it any further.

Gardner: And the failure mode than were basically lost records?

Moore: Lost records, couldn't find data.

Gardner: You couldn't find the data. It had been staged and somehow the pointers got screwed up and...

Moore: Yeah, exactly right.

Gardner: That's sort of interesting if it takes a mainframe to run HSM, how're you going to do it on a, whatever computing power you had in that control unit, you almost need a mainframe power to get things done. And then if you don't have that power you sort of play-- you're sort of forced to do a lot of tasks in microcode that are done for you and those are where your bugs get developed. It's a big job.

Moore: It's a big job, it was a big job.

Gardner: Okay, so VSS-- the other big thing that again was part of this proliferation of products was an optical program. What would you like to tell us about the optical program?

Moore: Well, the timeframe in the 1980's lent itself to people getting very excited about optics as a storage technology. And Storage Tech didn't ignore that, in fact the belief was that they could build a very high capacity optical disk product, they built a complete facility in Longmont, 10 miles north of their headquarters in Louisville to do this. In fact, Juan Rodriguez headed that up for Storage Tech so you probably got all you wanted to hear about optical already, but the...

Gardner: Actually, Juan and I haven't had the chance to discuss optical yet, so this is in certain way is a learning experience for me. Make me ask better questions.

Moore: So a great, huge facility was built in Longmont to do optical disks. The program was funded; it was called the 7600 optical storage subsystem. The product came out; it had eight storage directors per control unit. Data rate was one and half to three megabytes per second, it was WORM technology which was write once read many, so at that time, WORM-- remember I'm looking at this from a system and marketing perspective, all data didn't apply. Some data I needed to updated. In fact a lot of data back then with databases from IMS databases and CICS and things that were used on the mainframe; you did a lot of updates. If it's write once, it didn't work. So it was kind of relegated off initially to the archival space. And my belief is a lot of people at Storage Tech ignored that comment or didn't understand it. They thought here we have a great big thing that's going to work like a disk. But WORM was very different. A lot of data didn't work on WORM. It's got to be refreshed. Today that's different. We have an incredible amount of archival data that never changes in the world and WORM has much greater appeal than it did at that time.

But that product came out, the WORM disks were four gigabytes on one side and it was a 14 inch disk platter for optical disks and all that media was built up here in Longmont at one of the more fantastic manufacturing assembly lines and media production lines you ever saw. I can't remember how long the line was but it was 20, 30 feet long to put a piece of media out of that.

The seek time was 61.9 milliseconds, average seek time. That compared to 25 to 30 millisecond average seat times for rotating disks.

And the media life was projected to be 10 years when the product was announced.

Now on the heels of VSS here's optical disk, another interesting technology. Customers were really starting to like this one. And again, ultimately to make a long story short, the program was terminated. Optical disks never achieved its goals. The media never stabilized. The media was a huge problem to make repeatable media, make it-- repeat the process of media manufacturing and make it the same. And in the-- several years were spent funding this program, I can't remember how many were spent, but optical was funded and funded and all along, while optical disk was sitting here, you know, at four gigabytes per disk, magnetic disk was getting bigger, 30, 40 percent a year. And pretty soon magnetic disks were getting bigger than optical disks. Plus they could do reads and writes to them and they were faster and they were more reliable. And the bottom line as time went on, not just for this product but for the entire optical disk industry, optical disk fell way behind all magnetic technologies in price, performance, capacity, reliability, and throughput. Oh for five. It couldn't keep up. And today, optical disk is no longer as if it ever was a data center technology.

Gardner: You were talking about optical's inability to compete with magnetics. Who were some of the customers? The product actually did ship though right?

Moore: There were a couple of shipments out there and I do not remember the customer names on that. You know, the program really never got out of the batter's box. Some got to first base, some never got

out of the batter's box. But by the time this product was finally cancelled, and I can remember a meeting where Juan Rodriguez addressed all the employees and said it was going to be shut down, magnetic disk was running away from optical technologies. Today we have nobody really building optical disks for data centers anymore; the last company, Plasmon, basically is closing their doors. Optical disk does thrive however in your car and in your home. It's an entertainment technology. It's CD, DVD, it's Blu-ray, you know Blu-ray disk are 23 gigs, you know, the biggest disk drive today is two terabytes. Just to show you how far and advanced magnetics have blown away optical, so this was an adjunct product at the time. If we would have known then that the rate of improvement for our optical technology would be so slow in terms of capacity and performance compared to magnetics, I don't think the program would have ever gotten funded. But the fundamental issue was they could not shrink the size of that spot for a bit nearly as fast as the magnetic industry was improving aerial density. In the aerial density in the magnetic industry was a runaway train up until just about four years ago when it had to shift into perpendicular recording from longitudinal recording. And then that slowed it down a little bit for a while and it took the numbers down from 60 percent a year down to 35 to 40 a year. But it ultimately blew optical way out of the data center.

Gardner: In retrospect it was not a good technical idea, but when Juan told the team that they were shutting it down -- what was the impetus for that decision? It had gotten further than VSS right, it was actually a product, it was out of beta and but then it was cancelled. Was it the financial situation of the company or a realization of the product wasn't a good idea?

Moore: Huge financial pressure's always in the storage industry because what-- and that's a key one. In the storage industry unlike healthcare, unlike energy, where those two industries raise the price of what they sell 20 to 30 points a year, in the storage industry you lower the price of what you sell 20 to 30 points, where they're totally different economic models. By the way, none of those three are going to succeed. Two of them drive away all of the customers because they can't afford the service, the other one drives away the vendors because they can't make any money. In 1980 we had over 60 companies in the world building HDAs or disk drives, today we have six going to five left. And massive consolidation, there's not enough money, so the financial pressure was clearly there. But in my view, the bottom line was that the media process was not repeatable. They could not make a repeatable, high quality storage media and then they did not see a learning curve to take that media up and get it on the same path that it had to have to sell.
[1:00:00, 18:57/27:05]
It had to improve faster on aerial density than magnetics did to keep its capacity bigger and it could not do that. It fell way behind magnetics. It was, you know, optical disk was improving, aerial density at five to seven percent a year then, where magnetics were running at 60. And it could not keep up. So those issues, those combined, it just wasn't one thing.

Gardner: But did it have to compete with magnetic disk or was it really competing with magnetic tape? Well that's not moving quite as fast as magnetic disk.

Moore: Well magnetic tape at that time was not moving as fast as magnetic disk, so and there was no WORM tape so you could argue that it would take away a lot of the tape market.

Gardner: That was the argument.

Moore: But it didn't keep up with tape either. And by the way, in the year 2002 the capacity of a tape cartridge finally exceeded for the first time ever in the history of the storage industry, tape exceeded the capacity of a disk drive.

Gardner: Really?

Moore: For the entire history of the disk industry, if you wanted to back a full disk up on a cartridge or tape you had to use multiple pieces of media. Beginning in 2002, you could back up multiple disks on a single piece of tape because the capacity of tape-- and it's in my books here by the way well documented, had finally exceeded the capacity of disks. Today for example the biggest disk drive announced is two terabytes; the biggest tape drive can store two and a half terabytes with two to one compression. So, and by the way neither of those are shipping, but in native-- I could go back to native capacity same differences. So tape now has the higher capacity than disk and frankly that saved the tape industry in 2002 from the huge marketing perception that disk was going to become cheaper than tape. That did not happen on a per gigabyte purchase basis.

Gardner: Now VSS was actually dropped I believe before, well before the financial troubles in...

Moore: Yes. That was in 1982 type, '82, '83 cancellation.

END OF TAPE 2 / BEGINNING OF TAPE 3

[1:02:55, 21:52/27:05] in …3.VOB
Moore: What I'd like to conclude a little bit about tape and optical, I think the feeling at the time was that if optical disk would have been successful as a sequential access and random access product that was archival that it would have put a major dent into the tape industry at that time was growing very big. Customers were getting hundreds of thousands of tapes. In fact, I had been in four or five accounts in my life that had a million pieces of tape in there on the floor. And people wanted a high-capacity solution to get rid of a lot of that data that was not going to be modified. It turned out not to be optical disk, and the tape industry continued then to move along and kind of wait for the next big solution, because it was just getting too difficult to mount tapes.

So after the optical program was canceled, of course, you know, StorageTek was in Chapter 11. Financial problems from failed adventures like Optical Disk and processors, CMOS processors, major disk reliability problems. Those three things contributed heavily to the negative cash flow of the company and filing for Chapter 11 on that. Jesse Aweida left the company quickly. Company looked around for a few months, had an interim CEO name of Tom Wands, W-A-N-D-S, from Sears. He sat on the Board, and on Halloween of 1984 the company had filed Chapter 11. I remember a lot of people that I had in marketing at the time then, I was now the Director of Marketing at StorageTek, and I had all the disk product line. Another guy named Joe Beal, had all the tape product line and printers. He and I had disks and SEs. And he and I were kind of the only people left there in the company. Every day at lunch large groups of people would go out and talk about how bad things were. They wanted to leave. They didn't like the company. The ship's going to go down. And it was a very negative atmosphere, in and around, and

then right after Chapter 11. A lot of people were laid off at the time. So I remember going back to my hometown in Missouri at Christmas.

Gardner: This is now Christmas, ’84.

Moore: Christmas, ’84. And over the holidays, I got a call from Joe Beal, and Joe said, “Hey, everybody’s gone here. We’ve had mass exodus of the marketing talent, a people here.” There had been some people that wanted to stay just to live in Colorado. They would take no pay increases. They didn’t care. They wanted to stay here. They loved the state and the location so much. And Joe Beal said to me on the phone, I’ll never forget this. He said, you know, “We can build a marketing program with the people that are left here. Would you be willing to go with me into Tom Wands and present how we carve this up? So we carved up marketing between the two of us. Took it to the interim Chairman, Tom Wands, and Wands looked it over, and the next day said, “You know what? You guys got it. You both are in charge of all the marketing and go do it.” And here, when everybody was leaving, and he and I stayed during Chapter 11. This tremendous adversity turned out to be one of life’s great lessons for me. And that was that adversity breeds the greatest opportunity. And I came out of that as Vice President of Marketing, as did Joe. In the company within a few months, a brand new Chairman came in named Ryal Poppa, who was a turnaround expert. My hope was that as the CEO that he would be interested in more than just finances. That he’d be interested in the storage industry and what made it tick, and it turned out he was.

Gardner: Okay, before we get into Poppa, we probably should talk about the 8380, because I think its troubles may have contributed to the financial woes at that time.

Moore: The 8650 was what did that. We had 32 percent of the worldwide disk market with the 8650, and that started coming back in droves.

Gardner: Right, but then IBM started shipping the 3380, which probably caused some price cuts.

Moore: Oh, yeah! Oh, yeah, it snowballs from there. ’Cause a lot of your R&D money now went to product engineering to fix the older products, rather than fund the new ones, and the new ones got late.

Gardner: And IBM had announced the 3480 in that time period, early ’84, the 3480, so that put the cash cow business at some challenge. So then a lot of things were happening in ’84.

Moore: We’d hit from every angle.

Gardner: It was not a good year.

Moore: Hit from every angle.

Gardner: But also, I'm very interested in this plan you presented to the interim chairman. Do you have a copy?

Moore: No, we don't have a copy. It was presented verbally. And things were done real quick back then. In Chapter 11, they're making quick decisions, and there was no pound of paper to go with everything in there.

[1:08:08, 27:05/27:05] end of …3.VOB
Gardner: So I'm going to ask you to recreate that verbal presentation as much as you can. But let's step back a little bit in time, and talk about the 3380 class of product, the 8380. I believe it was announced about March 1983.

Moore: Okay, the 8380 product was StorageTek's answer to IBM's long-fabled program called Whitney, which was their 3380 project. And like other products at StorageTek, the company was always looking for differentiation, to have something better than that. And of course, now there were drains on the funds of the company to fund the failures of the 8650. Optical disk programs were taking money without offsetting revenue. And there was a big drain on cash coming from a lot of places here that were not producing revenue. So the 8380 made it in the market, and one of the first things that the company did, of course, was to go out and go right after the IBM base, so the newly installed 3380s. Fortunately, StorageTek had the cash control unit to go with it. And about a year later announced a version of the 8380 called the 8380R, which stacked twice as many HDAs in the same footprint. And this allowed people then to have a much higher storage density, double that of the IBM device, and this was appealing, because it had a smaller footprint. You know, it saved on the data center expenses and so forth. And people kind of liked that, and the company was able to go on and spin on an 8380 store. The 8380 never got the 32 percent market share of the 3380 business like the 8350 did for the 3350 compatible market at IBM. Market share never got that high, but the 8380, in my view, did okay. And it survived with a lot of its funding being pulled off to support other things that were headed to life support at that time. And it became the mainstay architecture for StorageTek with a variety of enhancements. The 8380R version, and there was one called the RQ that had an actuator-level buffer, which was a cache buffer built into each actuator within the system to offer even higher performance. That was not very successful. People didn't buy off on the additional cost on that at all. But the actuator level buffer was another very innovative concept. And some people had that with that. It's kind of like IBMs fixed-head disk on the 3350. You know, they announced that, but it wasn't very reliable, people didn't buy it, and it was killed off in time. So the 8380 was the mainstay product line really from probably '83 up through the rest of the 1980 decade for StorageTek; its base version, the R-version and then the double density version. All those enhanced that product line. Cash came with it, an actuary level buffer, and it turned out to be, feature-wise, a pretty rich family of products. But it was living off the legacy, you know, that StorageTek built an unreliable 8650. Would this be as unreliable? It was more reliable than the 8650.

Gardner: As it turns out, I believe the first 8380 shipped in March of '83. And that's about ten months after IBM's actual shipment, which had been delayed because of technical problems. And I guess IBM cut the price in November '84 on the whole 3380 product line, and that caused some of the financial difficulties, 'cause STK responded with corresponding price changes. Any interesting problems with the 8380, and I guess the double-density version during its lifetime that you recall?

Moore: Well, you mentioned pricing on that. In the storage industry, price cuts occur all the time. And price cuts affect your list price. And as we said earlier, the deals were made way off the list price, so you could make, in the storage industry, a lot of commentary around press releases for price cuts. But basically they had very little to do with the ASP, the "actual selling price" of the products that went on. Deals were cut, people would do anything. They'd go below margin if they had to, to see the big account. And the industry paid for that because the storage industry's number of players and profitability shrank on that. The 8380 products, like other disk products, you know, there were disk crashes; along the way there were other failures. But the company suffered because it could not get as much penetration in the market as it did with the 8650. It was more of the stay-even, catch-up technology with IBM. Where the 8650, when it had the double-density 3350 that IBM never did, you had a clear-cut advantage with twice the capacity per actuator. You had a price advantage on price per gigabyte. You had an advantage on data per square foot. And it was tougher to stay ahead of that. The 8380R did help that. But StorageTek then in the marketplace was, there were some-- people had questions about-- around their ability to build a quality disk program.

Gardner: Coming off the problems...

Moore: Coming off the past. The history had left some scars, and a lot of people-- and I spent a lot of my time working with customers to help convince them that we were back and ready to build a high-quality disk program. Tough job!

Gardner: IBM's failure to do a double-density 3350 created an opportunity for everybody in the marketplace. And they never made that mistake again. The 3380 had three mid-life kickers. That's, you know, they learned from that experience, and they probably took the market opportunity away from StorageTek and the rest of the market until the rules were changed by EMC, but that's beyond this discussion. So I'm very interested in the interim chairman?

Moore: His name was Tom Wands. W-A-N-D-S. He was on the Board from Sears, elderly fellow.

Gardner: So you and Joe...

Moore: Joe and I went in there and we said, "Here's what we'd like to do with marketing. There's nobody here. Nobody's running it. The people are running around the halls, not knowing what to do." Joe was a tape guy his whole life. I was more on the disk side. I was a storage guy in SE. And we just broke everything down in marketing, based on our own personal backgrounds, said, "Here's what we'd like to do." And the next day, Tom Wands says, "You know what? You guys got it."

Gardner: So what did you want to do?

Moore: Well, first of all, we wanted to begin to re-message the fact that StorageTek was still in the tape industry. We wanted to re-message that we're still in the disk industry. We wanted to rebuild the customer confidence in the company, so we created a user's group called the StorageTek forum, which was to come back and bring people together, you know, to realize that we're actually still alive. And Joe and I

were both big on marketing and messaging. So it was an outward-bound thing. We set up educational programs to bring people into the company to let them know what was still going on back here. Because communication had shut off in the downturn. And shortly after that, to give us a big boost was the hiring of the new CEO Ryal Poppa. And he brought Joe and I, and probably a dozen other people down to the boardroom one night about 6:00. Everybody's thinking about going home, and he says, "Guess what, guys? We got a meeting." And I remember that boardroom being full of the people that he designated as the key people to bring the company forward, and rebuild it. And he said, "I want to see your plans, and everything on what we're going to do here to get this thing back." He said, "We will get the company back." And he was a turnaround artist in his past. And we didn't know him. We'd never heard of Ryal before.

Gardner: The Ayatollah? He was called "The Ayatollah" for his past turnaround experiences. And that was before the Ayatollah had the connotation it has today, but maybe it is the same connation. What was it like?

Moore: Yeah, yeah, working for Ryal was outstanding. It was one of the greatest experiences; maybe the greatest of my professional career. I remember the first one-on-one meeting with him, I thought, "Well, he's going to want to know about financials, and whatever," and his first question to me was, "Let's talk about HSM." And right then and there I knew he was into storage and how to optimize things, and we started building a storage strategy around that. And he was brilliant. He was tough. I don't know about "Ayatollah." Today that certainly has a poor connotation, to say the very least in global politics. But back then Ayatollah may have meant something different, and I was unaware of that. But he was a man who made decisions. And he looked at data, and then he would make a decision. He wouldn't spend months and months just studying stuff, and having people grind out endless surveys. He'd see enough data and say, "Okay, here's where we're going to go." Very customer savvy, spent a lot of time personally with customers understanding them. He was deep legally, financially, technically he'd been an IBM branch manager, and knew that game.

And he entitled us all to make a lot of decisions. So it worked because we were now empowered. We didn't have to go get approval for everything we were doing. Said, "I need you to go and build me a strategic plan. I need you to build me a tape plan for two years." And people became empowered pretty quickly. A few people didn't like that. You know, they wanted to be told what to do. But he wanted people who were leaders, not necessarily followers, to be part of that. I thought his leadership was outstanding. I never had a conversation with Ryal that I never came out of it feeling better and more excited about the company until the day he left.

And he truly brought back the ship headed to disaster. He had to fight off the creditors, to keep the tape library alive. I mean, optical was ultimately gone with the creditors. The disk business people, several wanted to shut down, because it was marginally profitable, based on all the rework of the disk failures. And he fought hard to make that happen. We had workout groups, met with the creditors constantly, to show them what we're doing. And it worked. And then we started going to Wall Street.

And then in 1987, we went out, in April I believe, and announced the 8380E, double-density 3380, to once again do what we did on the 8350, and jump ahead of IBM. And the 8380E was significant. The announcement was done at The Tavern on the Green in New York. Wild crowds there. We were about

ready to emerge, later on that year in August I believe, from Chapter 11. And the analysts were all over this at The Tavern on the Green with the 8380. “You’ve come back. You’ve built a new disk drive in Chapter 11. You have a tape library underway,” that they were dying to hear about that was coming along here.

And here as we began to exit Chapter 11 process, was a new disk family, the 8380E, and a tape library to come out of that. And those two products elevated the company back, and showed that 1) that we could jettison the things that were not profitable, and not working well; and, 2) had enough resources left to build two good products - the 8380E, and of course, the tape library, which was far better than just “good.” Turned out to be a huge success in the tape industry. And those things led a revival of interest back into StorageTek. Ryal brought on another guy named Steve Jerritts, who does live here in Boulder, as the president of the company. He brought in some top-flight people, and his senior team was intact to go forward. But he had a well-orchestrated series of events to rebuild the company. And we all wondered, “How did this happen? Would it just be willy-nilly? Or would it be...,” but he had one step at a time. Very carefully done, and he kept the people that he thought would really do it. I was delighted to be part of that. I never had a legal degree, but after going through three years of that, it felt like one. And here we were coming out of Chapter 11 in August of ’87 with some exciting things. A new disk product, a tape library that was looking extremely good, and the seeds had been planted for something at that time that nobody then knew anything about, called Iceberg, which was intended to revolutionalize the entire disk industry.

Gardner: So let’s back up to Poppa comes onboard -- he inherits what?

Moore: He inherits a disk product line that was in-- I don’t want to say “collapse,” but struggling after the 8650 reliability issues, which were widespread and heard around the world. He needs to bring the 8380 into the limelight.

Gardner: Right. Okay, Poppa joins in March 1985. So the 8380 is announced, and shipping...

Moore: But remember, you had 32 percent of the disk industry out there. What’s announced, and what’s in the field are very different.

Gardner: Of course.

Moore: And there was still this huge field population or products that were failing, despite anything.

The challenge was to get those people to convert to the 8380 family, and believe that the new one was going to be better than the old one. Because he had to preserve a base that was eroding and defecting and going elsewhere.

Gardner: So the challenge was to manage the base to transit to the new product.

Moore: With fewer people. A lot of sales people, marketing and engineering people left. So, he had fewer tools to work with. He had a bigger job.

Gardner: And the base tape business. It was a cash cow, and thanks to IBM's strategic errors, STK dominated the reel-to-reel tape business. But there's a new 3480 announced, which has to do something in that business.

Moore: Yeah, and of course, the day that's announced is not like, "Whoa, here's their tape cartridge." You know, for a year or more you're tracking. You have all sorts of eyes and ears in the industry. You know, and have a good idea what's coming about. So when that cartridge came out, the cartridge is what enabled a tape library to exist. Because Robot could pick a cartridge. But to pick around a reel of tape was virtually impossible. So there was already a concept to have a robot in mind that could pick out this cartridge. Whatever shape it was. It turned out to be five by four by one inch in size.

Gardner: So that's a good point. The concept, then, in the tape business was to respond with not a tape-- not a conventional tape drive. Tape control unit, the way the 3420 series was, but to go directly to a library with the cartridge, or was it the plan to do both a conventional tape drive, control unit, and a library. What did Poppa inherit from the current team, and what did he decide to do about it?

Moore: The main focus was on the tape drive to be compatible with IBM. The library was essentially a skunk-works over here that may or may not be funded. Proof of concept type of thing. It was not a mainstream part of any strategy at that time. And he had to be convinced of it, and then once he got convinced of it, he had to convince the credit committee. And you know, all the creditors and everybody that this was worth funding.

Gardner: And the library was intending to use the 3480 cartridge.

Moore: Yeah.

Gardner: Even though he may not have even seen it yet. It was just announced.

Moore: But you knew. You know, everybody had a sense, you know, it's like before the 3380 was announced for your-- you had a sense. Was it 46,000, 48,000 bytes or whatever. And the same thing here. So yeah, there was a mock-up library built to retrieve cartridges. 'Cause you could grab a cartridge. You couldn't grab a reel. So that was kind of betting on the come process.

Gardner: Now is it true that that project had to be hidden from Aweida? There were stories in the papers of afterwards when the libraries came out that Aweida was opposed to, and the product was actually done behind his back.

Moore: You know, I cannot verify that. Working there, I had no sense of that at all.

Gardner: But when you became VP of marketing after Aweida left?

Moore: Yeah, under Ryal. After Ryal came then he elevated us to VP of marketing, Joe Beal and I. No, the tape library, he took a look at the tape library. He talked about the tape business, the archive business, the long-term storage business. And with the help of his friends, I think, decided that a robotic library was in order. Remember, IBM had tried a robotic library on the MSS 3850, and I would argue it was not successful. CalComp had a library, Memorex had a library. None of them were very successful.

So, the problem that Ryal saw was this, which we saw, that I saw in marketing, data centers were overflowing with tapes now. People were on roller-skates in some places mounting tapes. I remember when they had a backpack, they put them in their hip. They had contests to see who could mount the fastest. “Here’s our fastest mounter this week, you win a free meal out this weekend.” It was unbelievable. One place we went said, “You can see the curvature of the earth at the end of our tape library. I can remember seeing a place with a million pieces of tape.

And had this library not been successful after some other marginally successful, at the very best, starts, I think the tape industry would have died. But this one made it. This one used circular robotics, it was reliable. By then, robotics had gotten much more reliable than before. So this thing took off. And Ryal believed. He says, “You know, the tape industry, our cash cow will die, because people can’t manage it anymore.” By the way, optical had died. Though a lot of people were still thinking optical would be the replacement for tape, it just couldn’t get out of the batter’s box. So now you had the tape industry under great fire. Too many tapes, too many mounting errors, lost tapes, people handling every piece of media. That was just not working in the marketplace. Particularly for the Fortune 1000 customers who were heavily relying on mainframe tape.

Gardner: So who was the marketing advocate that got the robotic library project going? Was it going even before the 3480, or did the 3480 sort of inspire that? You make the point that the cartridge is much more amenable to robotics than the open reel.

Moore: Yeah, there would have been no successful robotics for the open reel. So the cartridge was the catalyst that made a robot possible. Who made that take off at StorageTek? Joe Beal was certainly one guy who ran that tape business then at the time. He had a guy named Ron Korngiebel, <laughs> K-O-R-N-G-I-E-B-E-L, who was very dynamic, hard-nosed marketing guy that helped fight that thing through. Yeah, and they drove it through, I would say. And he had help from a lot of the rest of us in the storage business. I mean, I certainly believed heavily that robotics were the key. Circular robotics meant a finite amount of time to anything, where a rail, you could take forever, variable times. So you know, the circular robots had a great appeal over the traditional rail.

Gardner: Whose idea was the circular robots?

Moore: I don’t remember. I don’t remember, but that was a great idea, whoever wanted to go circular on that.

Gardner: Because he MSS was a rail, and you could, in fact, bolt more rails on. I think three or five.

Moore: Santa Fe on the side, you know? Several boxcars long. And Tim Studebaker may have been the main engineer on the robotic library. He still lives here in Boulder. He was a lead robotic architecture. Later went to work for Juan at Exabyte, by the way, on their libraries. But he was one of the pioneers on this. So, they had a great team of tape engineers. And both mechanical for the robot, and electrical on the drive side, and the tape engineering team was excellent. Now tape, remember, had to fund all the other programs that didn't make money in the company. So not only did they have to fund their own future, they had to fund everything else, 'cause that's where all the positive cash flow was coming was out of the tape business.

Gardner: So basically, the disk business was probably not profitable, because of the rework cost and the price cuts, cutting.

Moore: Right.

Gardner: But pretty clear what to do there. We talked about the transit to the new product, preserve the customer base. Not so clear what to do in tape. Have to do a 3480 equivalent, because that's your cash cow business. But we have this new library thing, which is not so clear, but Poppa becomes an advocate of that. Optical is shut down. VSS is already gone. I guess Aweida had already shut down the printer business and telecommunications business.

Moore: And a CPU business.

Gardner: Had Aweida already cut the company down to where Poppa could then grow it forward?

Moore: That's a good question, and I'd have to do a little research on that. I think there were probably a number of OEM projects that were still hanging around, that weren't bearing fruit that had to be trimmed. How far you'd continue to fund the old 8650 family, and shifting funds to newer things like that. There has been some cutting done at StorageTek. But it was not all over. When Ryal came in, there was still more cuts to do with wastage and facilities and operations. You know, and you had to take-- in fact, Steve Jerritts his Chief Operating Officer and President-- was a guy that really came in and cut even more additional expense out of the operation side. Manufacturing, building your own media, you know, "Whoo! Let's get that clean room out of here. Let somebody do that." Those kinds of decisions still had to be made on that.

And the tape thing, of course was, "We're going to build this cartridge tape. And by the way, we can save the tape industry," though that wasn't any pubic message, "If automation is successful." 'Cause it would take people out of the loop of placing data on the channel. And this was a significant event in the history of the storage industry. A successful robotic architecture. Now human error was pretty well gone from this whole thing. You know, 6,000 cartridges, roughly, in a single library. It solved massive labor problems for mainframe systems. It did not initially connect to the non-mainframe world. So you had all the Univac machines and all these other computers that were-- that did have tape in much smaller degree than a

mainframe, but they had no real access to a tape library. And their need wasn't as great. The tape world was really centered around IBM mainframes that used tape, you know, with a voracious appetite on that. And that library came out and, you know, from then on, 14,000 units later, it was one tremendous story.

END OF TAPE 3 / BEGINNING OF TAPE 4

[0:00:00, 00:00/26:50], [from beginning of DVD2, from beginning of …1.VOB]
[0:00:10, 00:10/26:50] begin…2.VOB

Moore: We look at the tape library business as probably saving the tape industry from extinction. As I mentioned, a lot of people had so many tapes to mount, so much confusion, the bigger the data center the more tapes were there and the more problems that occurred, and the dissatisfaction, by the mid 1980s, with tape was huge. And StorageTek, with a lot of help from the Chairman, saved the tape library project, made it a full blown, full funded project in the face of the creditors and Chapter 11.

And in 1988 at ACS, the Automatic Cartridge System was announced. And this system came out as a circular robot. The robot went around in a, literally in a circle. It could hold up to 6,000 cartridges for a total capacity of 1.2 terabytes in that library. Consider that today that one tape cartridge itself can hold 2.4 terabytes today. So in just that 20 year period of time the capacity of a tape cartridge has gone up a thousand times. And when people thing the tape industry's dead or hasn't done anything they haven't followed it very closely. So one cartridge today holds as much as the entire library did when it was first introduced 20 years go.

And the robotic mount time was good. It could 90 exchanges or mounts per hour. Back then for a person the average tape mount time ranged from 30 seconds to 90 seconds, so this thing was doing 90 exchanges per hour in that robot. That was not bad. And people had multiple silos, and they tried to automate as much of their data as they could that had previously been on tape.

Remember, the library only worked for the cartridge. There was no way to make it work very effectively where a robot could go and get a round reel. So robotics took off, saved the tape industry. StorageTek ultimately continued to enhance the library family. When you think of a library you need to break it down a little bit. You have the basic library, but over the life of that library the drives come and go. They're replaced by higher capacity, higher performance technologies. The media comes and goes because the form factor did not change in a library. One of the great things about a library was, unlike disk, the form factor didn't keep changing each time something new came out. It was always about 5x4x1 inch and it would fit in that slot. So you could put new media, new drives, constantly upgrade that same footprint.

And here we are today 21 years later and a large number, thousands of the robotic tape silos, the round robotic libraries are still installed. This easy upgrade capability made people really like the product. So we didn't have to have all the tape mounters, all the mistakes. The reliability got a lot better with libraries.

StorageTek continued to enhance the family of drives that went against that. And one of the biggest challenges with StorageTek occurred in 1998 in an effort driven by the then general manger of

StorageTek's tape division, a guy named Gary Francis, and that was to break-away from being compatible with IBM's device image and media. StorageTek would make their own, have their own media, their own drives, their own capacity and it would be different than the plug compatible model, what they've used through the 9-track in the 3480 version. And this series of products called the Titanium Family, the T10000, was a brand new version of products, different capacities, different specs than IBM with that, and we really kind of lost, at that point, the two remaining players in the mainframe market, IBM and StorageTek from being interchangeable. Since they each had their own media there was no more media interchange, and that bridge was forever broken with that direction. It turned out to be very successful for both companies going forward. StorageTek was able to innovate and enhance much faster rather than waiting on IBM to introduce something and then reverse engineering, and IBM could do what they wanted to do as they always did. So that was a major shift that got marginal visibility in the industry, but the first time ever that cartridges and 9-track tapes and everything were not interchangeable between IBM and StorageTek in the mainframe space.

Gardner: Correct me if I'm wrong, I thought the 3480 was 18-track or 36-track.

Moore: You thought correctly. It was both.

Gardner: You said 9-track several times in this. I thought the 3480, or the ACS used 3480 cartridges and those were IBM compatible cartridges until 1998. Is that what I...

Moore: Right, and then StorageTek put their own media format in there.

Gardner: Was it still in physically the 3480 cartridge?

Moore: The physical package never changed. What's inside of it was different, number of tracks, number of read/write heads, all that stuff. They broke the reverse engineering mode and built their own drives to record data their own way.

Gardner: Okay, but still a physical 3480 cartridge the same, I think it's still half inch tape?

Moore: Yeah, still the same physical package.

Gardner: But now different recording formats, different media. The tape drive, of course, had changed.

Moore: It looked the same. The robot saw the same kind of package. The robot didn't know what was in there, but it put it on a drive that could read it, but you couldn't read a StorageTek cartridge on an IBM drive any more or vice-versa. They had broken that mold of interchange.

Gardner: Now, were there a family of StorageTek drives and conventional controllers to go along with the ACS or how did StorageTek continue in the conventional tape drive control unit attached to IBM channel business, if at all?

Moore: So when the library came out all the StorageTek drives could either be automated drives on a library or they could be standalone as in the past. At that time a lot of people used cartridge loaders, stacker loaders that they called them, and they could just use a standalone tape drive and have people mount it or use the loaders, or they could put those very same drives on a library and automate it as the shop got bigger. The bigger the shop in terms of tape utilization the more likely you were to automate. Just think, the entry level point at that time was 6,000 cartridges, if you only had a hundred you probably weren't a candidate. So then that actually made that. And then StorageTek brought out, continued to enhance versions of tape all the way up through. And then they were compatible with IBM up until 1998 and then they broke the compatibility mode.

Gardner: So but then physically the same drive could be mounted in the ACS and also that's surprising because I think of an ACS drive as a rack mounted small type device and I think of a standalone drive as something that's about the size of a refrigerator. I guess I'm confused.

Moore: That was really good news. Your confusion is over because you could do that and you didn't have to have a separate set of drives that were library only drives. You could have a data center. You could have 8-16-32-drives and when you got big enough to automate you could put those drives up against the library. They would take away from library capacity because they would take away wall space where some slots were like that, but they bolted right up to the library.

Gardner: As opposed to a redesigned rack mounted drive that slid into a rack?

Moore: Yeah, that's double the R&D and engineering cost to do that. That was a poor economic path to go, so they decided to let people be able to upgrade their existing tape infrastructure to an automated one whenever the time came.

Gardner: Using the same actual drive?

Moore: I know you don't want to believe you, but it's true.

Gardner: No, no. I believe you. I'm repeating just to make sure the record is clear, because I have a vision of libraries like other people have made or even the idea of MSS where the tape drive did not take an entire quadrant or segment of the library, but you put multiple tape drives in a single segment. That's sort of the vision I had, had I looked carefully at the ACS.

Moore: This was to allow the customers ease of upgrade, not have to swap out gear, resell it and/or take it off the lease. They loved that strategy. The interchange and flexibility on these libraries with drives and media was excellent. The thing you couldn't do on a library was use your scratch loader. That didn't

make sense any more. The tape library had scratch tapes within them to go and mount in there, so that 10-deep, 8-deep scratch loader architecture that people had seen was not applicable to a library. If you had that and you wanted to move that drive over to a library you couldn't do it.

Gardner: And then there were conventional control units to attach these drives to a channel if you weren't going to be in the library business, or you could have both.
[0:10:00, 10:00]
You could have a channel attached, control attached drives, and drives in the library which were also channel attached through a control unit in the library.

Moore: All the drives had their same channel attachment and hardware, but the library architecture had two things. It had a control path for all the commands and it had the standard data path that existed on all the tape drives. So all the library commands robotic positioning, go to cell eight, row three, column two, all that came down the control path, and by breaking the connection with a control path and a data path StorageTek then added its own library LMU, Library Management Unit and host software to make it work. This did require a sizeable host software component on the MVS system. Later on when StorageTek attached that library to non-mainframe platforms what they did is then they moved that software into another device and that allowed non-mainframe operating systems to begin to use the library. Unisys was a great example of that. They were one of the first big users that had enough tape business. Even though their tape pools were much smaller than mainframes they had enough to warrant a library or two, and that allowed the library to connect to other platforms. Today those libraries connect to mainframes, UNIX, even a Windows system, but Windows doesn't do much with tape, but to UNIX platforms mainly, so.

Gardner: So the library manager was a separate device addressed by the system and that's how the command and control was sent out.

Moore: That was a separate device.

Gardner: Essentially the mount/dismount commands.

Moore: All in the control path over here to make it work.

Gardner: Separate from it and then...

Moore: Yeah, once you had the drive with the cartridge mounted together then it was standard data transfer.

Gardner: And in fact that same control unit could have a drive in the library and a drive standalone.

Moore: Uh-huh.

Gardner: I mean up to the number of drives and control units.

Moore: You could. I don't know why, but you could do that, yeah.

Gardner: Well, if you had some-- I don't know why you'd want to do that either.

Moore: Yeah, well the library got better. So we had one library that had up to 6,000 slots in it. You could bolt together, ultimately on the first round, 16 of those into a whole automated cartridge subsystem. So you've got 96,000 cartridges together. And the robotic hand was really interesting. The robotic arm could take a cartridge, and if all the drives were busy in this library, say you had six drives attached to the library, it could go through a pass through port. The hand in one of the adjacent robots could take that cartridge and put it on an available drive in that library. So that was called the pass through operation. It took more time, but it allowed you to satisfy and better optimize the utilization of all the tape drives.

Gardner: So was the 6,000 cartridge capability with two drives and one pass through? It's an octagon, a multi-sided device. Each side could be a function, could be an array of cartridges, or it could be a tape drive, or it could be a pass through port. Maybe the pass through port was part of an array of cartridges. So when you had 6,000 cartridges could you also have two drives?

Moore: On that library?

Gardner: Yeah, on that library.

Moore: I think you'd have two up to 16 drives, your choice.

Gardner: Well, then that's 16 sides around the device.

Moore: But you don't need 16 sides in the first place.

Gardner: Well, you need a side per drive don't you?

Moore: No, these are multiple drives in each frame.

Gardner: So the drives could be stacked one on top.

Moore: Yeah. Well, the drives came-- you had at least two drives in a box, so yeah you could just bolt up four boxes on there and get eight drives. You could put 16 drives up against the ACS. It was really highly active. There was room to do that. I don't know too many people that did that many on one actuator, because that actuator would have to really be busy to feed 16 concurrent tape drives, but.

Gardner: Did we do my usual question about key players on this architecture both the physical, mechanical?

Moore: On the tape library, we did. The library the key guy was a guy named Tim Studebaker, so we got him on there. Joe Beal in marketing was the program manager for that. I mentioned a guy named Ron Congable who led tape marketing. He was a highly aggressive marketing guy and very visible in tape, and really, really fought the StorageTek cause well. So the library, though, ultimately didn't need the kind of ultimate convincing that the disk products had a decade earlier. The library worked. People knew it. The word spread. And pretty soon IBM, who when StorageTek announced the library in 1988 IBM said, "We will never ever build a tape library." Four years later they were building two libraries. Very different than their statement four years ago, "We said we'd never do it and now we're doing it." So IBM got into the library business. Other people got in the library business and automation saved the much broader tape industry. Everybody realized that robots were better than people for placing data on a channel.

Gardner: Libraries existed before the ACS. Apparently IBM decided when they saw the ACS it wasn't a good idea, but other people thought it was a good idea. What was it about StorageTek's ACS that caused it to succeed and then StorageTek to really dominate the library business through the '90s?

Moore: The library business began long before StorageTek. CalComp had a long rail library, not very reliable. Control Data had a long rail library, not very reliable. Neither company existed in the storage industry by the mid 1980s. They were not players. IBM had developed a 3850 library on a rail and I would argue that it wasn't successful financially, so they killed off that program too. So that left the tape industry without any automation. Other than IBM the other library companies had no big install base of tape. They were not known as a tape provider. StorageTek had a huge install base of tape. StorageTek always ran neck-in-neck with IBM in market share. So on that big base was their great gift here to introduce a tape library and let people begin to take that successful 4480 compatible family of cartridges and automate it. So their base was huge. They were 40 percent, 45 percent of the market. The other players in the past didn't have anything like that except IBM and IBM steadfastly said, again, "We're not ever going to do a tape library." Later on they did two, but they're install base enabled. StorageTek was a reputable well established tape player.

Gardner: You were also first with the cartridge library, right? I mean, you were the first to jump onto the 3480 cartridge into a library which contributed to the reliability because the cartridge is much easier to load than a reel, but also perhaps the rotary actuator, or just a good design of a reliable actuator. Was it the rotary actuator or just the fact that it was reliable?

Moore: It was four or five years after the StorageTek library before others started doing it. I mean, the non-mainframe world was not a heavy user of tape like the mainframe world. IBM wasn't going to do it. That left StorageTek as the only mainframe provider. But the non-mainframe world got bigger in terms of tape. They saw that automation was good. They didn't build 6,000 cartridge libraries. They built 100 slot, 100 cartridge, 200, 500 cartridge libraries, maybe a thousand, but they weren't as big. They served smaller tape environments. But sooner or later anybody in the tape industry understood that for them to survive they needed a library. You just couldn't afford to have people hanging tapes and filing them away all the time. From security problems on down the line that didn't work, so StorageTek made libraries popular. At one point in time I'd counted 14 companies delivering some from of tape library. That does

not count scratch loaders with that. Today there are far fewer tape companies in the game, but automation is the key thing. In 1998, just to go back a bit, both StorageTek and IBM introduced and implemented the integrated virtual tape library. They put a disk array in front of their tape library integrated in the same physical hardware, and that disk array was a buffer. That disk looked like several virtual tape drives to the operating system and a lot of the reads and writes could come off of that disk and not have to go to backend tape. And the integrated VTL became a dominate concept in the mainframe world. And that has been the last major architectural change that we have seen in mainframe tape since then.
[0:20:00, 20:00]
Since then we've had enhancements on that, but the disk array embedded in the tape library took 60 to 70 percent of the IO activity off of the backend tape which made the tape more reliable because it had less wear and tear and saved that IO activity to and from disk. Backup and recovery was a huge beneficiary of the integrated VTL concept.

Gardner: Anything else you'd like to tell us about tape libraries?

Moore: Well, in the history of StorageTek this was probably the product that the company was best known for. It was reliable. It worked well. It solved huge needs. The price was certainly in order with what people want. You could interchange drives and media without removing the library. Where most products have a five year or less lifetime, this thing ran 15 plus years as a library. It was growable with pass through ports and multiple groups of 16 libraries all connected together. In fact you could do 16 of those groups if you really wanted to go huge and get into the gigantic scientific market space. So that was really a major flexible architecture that turned out to be very well thought out. I remember StorageTek's first users' group in 1988. We had the first users of the tape library get up and talk about their user experience, and their biggest concern was what to do with all the packing material. They had never seen so much cardboard and packing material from all these panels and everything, and recycling wasn't a big issue then, so that was the huge issue. But this was a homerun for the tape industry and maybe the biggest one in the history of the industry.

Gardner: I think the next most interesting product might be Iceberg, but let's just take a segue into XL/Datacomp and Alpine because I understand this is really two different companies. I mean, the mainframe world that you and most of StorageTek was in, was really different than XL/Comp and Alpine, which is the product I'm interested in. So why don't you elaborate on how XL/Comp and StorageTek related?

Moore: Well, StorageTek was looking for new growth opportunities in the late 1980s and the then popular IBM AS400 series looked like a great place to go. So the company thought, "Well, we'll take a look at this market, IBM's behind it." It was the midrange business machine of choice and ultimately StorageTek purchased a company based in Chicago called XL/Datacomp who was a huge reseller to the AS400 marketplace.

Gardner: And they were pretty much separate from the folks in Louisville?

Moore: They were separate. The acquisition, I would say, was difficult to integrate into the people in Louisville. It was people in Chicago, people here. They kept their culture. They kept their same values.

There were some products developed to help them do better in that market. The product you mentioned, the Alpine product, I believe you said, right?

Gardner: Yeah, Alpine, it was OEM'ed from array technology.

Moore: Right. The Alpine product was being developed in Colorado to go in that market, but right after StorageTek acquired the AS400 IBM closed off the IO adapter to allow plug compatible machines to easily attach. They had a black box architecture on the AS400. It was a single level storage system and you never knew where your data was in the AS400, even today. It's now called the I Series. You never knew where it was, and that made it very difficult almost after the acquisition for XL/Datacomp to sell anything in that market, because with the new IO adapters on the AS400 you couldn't plug in. You had reverse engineer and do a whole lot of stuff, and that lead to that, "Build our own disk array now. We just can't go plug into that market," and it never worked. XL withered away and over 98 percent of all the storage connected to an AS400 comes from IBM. There really is no plug compatible market. So that whole acquisition essentially vaporized over a period of about four to five years, because the market you just couldn't get on that.
[0:25:00, 25:00]
I mean, even today people try to put a few tape drives or a disk drive on the AS400, very, very little business there. It's difficult to do with hardly any ROI. So that was a strike and an expensive one. The people were let go and it affected developments like array technology that obviously even if you could have built it you couldn't have sold it.

Gardner: We then turn to the Iceberg saga. What would you like the world to know about Iceberg?

Moore: Well, go back now to 1988 and '89 and what's going on here. And the company's got a successful library launch. They have come out of Chapter 11. There's momentum and interest back in the company. Revenues are growing from a low of $636 million, I believe in 1987, and there's momentum back. And the company, and the whole disk industry is coming off of a history of marginal to very bad reliability depending on who you were in the game, but disks were unreliable. They failed and that data was un-accessible. So there was a concept that had been disclosed by Cal Berkeley called the RAID architecture and it was then stood for Redundant Arrays of Inexpensive Disks, but later became Redundant Arrays of Independent Disks, and whatever else you wanted to call it. And RAID allowed a new way to lay data down on the disk, and it had data written on one disk and RAID 1 would mirror disk, so whatever I wrote on one disk I wrote on another one.
[0:26:50. 26:50] end of …1.VOB
Another version of RAID would have parity with it so if the disk failed I could reconstruct the record that was lost on a disk with parity and the remaining data. And there were even more advance versions of RAID like RAID 6, which was called log structured files which allowed you to do all sorts of updates. It had dual redundant rotating parity. StorageTek's history in disk reliability was tough. It was one of the primary contributors to Chapter 11. And at that time Chairman Ryal Poppa said, "We need to find a way to change the rules of the disk industry."

Gardner: Let me just interrupt. Actually, StorageTek announced that they were going to do Iceberg in very early '89 and I think that's even before the RAID paper.

Moore: The RAID papers were before that. Katz and Patterson had published the RAID specs at Berkeley before that. Nobody had built it.

Gardner: Eighty-eight actually was the RAID paper. January '89 StorageTek announced Iceberg without a date, sold out. Well, then it formally announced it in '92 three years later, and went through some problems and I think ultimately shipped in '94?

Moore: Yeah. Iceberg was announced on January 29th, 1992. For two years before that thousands of customers were under nondisclosure presentations to get their feedback all over the world, literally, on Iceberg. "What do you think? How will this work? Will this solve your needs?" So there was a huge preannouncement market research campaign going, biggest I've ever seen by far, and to do non-disclose everybody in the world essential meant it was public information when it was announced in '92.

Gardner: Yeah, it was published in January of '89. I mean, I'm looking at an article, 5 January 1989. It's the first public mention of Iceberg. Maybe somebody violated their nondisclosure.

Moore: Yeah, that's not an announcement, though. Public mention and an announcement are different. It was announced in New York at the Hudson Hotel in January 1992.

Gardner: Somebody disclosed the word Iceberg associated with an array system as early as 1989.

Moore: Well, StorageTek did that, and they did it at the Annual Gartner Group Storage Conference, because one of the things that chairmen were asked to do at Gartner Group was to walk about your future. Where are you going to go, and we have a program-- what was said is, "We have program code named Iceberg that's coming out." Then everybody said, "What's Iceberg," and the questions began, and then the nondisclosures began, and it was a way to get significant feedback from people.
[0:30:00, 3:10]
And pretty soon everybody knew what it was. I mean, there must have been five or six thousand nondisclosures done on Iceberg. And it was going to solve every problem that existed with count-key-data architecture. This was the coup de grâce of all disk storage systems. Still today many of the concepts imbedded in the original Iceberg have yet to appear.

Gardner: So what was the '89-'91 pre announcement concept? Could you elaborate on it?

Moore: Well, the preannouncement concept actually was the same as the announcement concept, but we had a lot of ratification and a few enhancements made by the customers and that, but Iceberg turned out to be a disk array. The initial prices were $1.3 to $3.6 million dollars per array. It had 400 gigabytes of native capacity, but what Iceberg did was lightly compress disk data, and this has the first disk subsystem and the last to ever do outboard compression of disk data. When you compress data you can't update in place, so that necessitated the architecture to move to what was called log structured files. It was also going to use RAID because the reliability of disks then were not very good regardless whose they were. They went from average to not good. So you'd used RAID on top of that so that when

a disk failed the file could still stay available, and then you could rebuild the failing disk on a spare disk in the subsystem.

Gardner: And in the early embodiments was it multiple forms of RAID or a particular form of RAID?

Moore: It was a particular form of RAID called RAID 6 Plus.

Gardner: Originally?

Moore: Originally. It was log structured files all the way. We don't update in place and that enabled people to buy 400 gigs of hardware and store 800 gigs of data. Previous disk drives you bought 400 gigs of hardware and you stored 300 gigs of data. And for the first time ever the paradigm had shifted where you could store more data than you had actual hardware capacity. That was pretty neat, and that gave a great economic benefit to clients because the systems engineers would run software tools to see what their compression ratios would be. So if I bought a 100 gigabyte Iceberg system I might be able to put 250 gigs of data out there. And you wanted to leave some room for growth and buffer space and things like that, but the economics became initially extremely favorable on that. Iceberg implemented a concept now that's used all over the non-mainframe world called snapshot copy. Snapshot copies took a picture of all the pointers of the data like a picture at a point in time, and you could freeze all those pointers and leave that data. And as the file changed all those elements would stay the same. So your file at 1:00 in the afternoon would always look like it did. And you could use that as your test file, a backup recovery file. It was a great way to do backup and recovery, to tape off a disk and freeze it at a point in time. Otherwise, if the file was still active you'd have to close the file, shut down the application and back it up. So you could essentially do things like hot backups at the time. And that had great appeal. The space efficiency had great appeal. RAID had great appeal because the reliability of disk was far from perfect going into the early 1990s. StorageTek was two years late with the product.

Gardner: At least.

Moore: Yeah, well they announced it in '92 and they shipped it in '94. And the two years did dramatic damage to StorageTek. People began to question dates. They began to question commitment, engineering capability, how can you announce a product and be that far off yet? Here we're making plans to install it now we've got to scrap them for two years. And the delay on Iceberg was really costly because IBM had a marginal disk product of their own out at the time. It wasn't doing real good. They had a product that was ho-hum, and here was one that was going to blow IBM away in the disk area if it did what it was supposed to do. Here was one-- there was nobody around that was a real good competitor disk other than Hitachi at the time. EMC was a small mainframe company and they built the solid state disk with a big disk backend, but this was the great opportunity for EMC. EMC saw StorageTek couldn't ship. They saw IBM in a very weak position in the marketplace with disk, and they decided to do something really good.
[0:35:00, 8:10]
They brought out their solid state disk and they put gigantic backends of disk storage on there and a large amount of cache. They had a great cache subsystem. You knew that as Symmetrix. And what they did, also at EMC, is they guaranteed performance. They wrote down a performance guarantee. "We'll bring this in for 30 days, 60 days whatever. We'll guarantee very IO." I don't know the exact-- 5 millisecond

response time, and people brought it in. It was fast because it has so much cache. It over cached. When you have so much cache almost every IO is a cache hit. So they did that, and they took off, and they never looked back.

Gardner: Yeah, they actually shipped that product in 1990.

Moore: The small backend, but the big one came as the other vendors began to open up. They started to open it up and their market share really exploded.

Gardner: Their first product with one gigabyte disk drives was a 24 gigabyte, had 24 drives which could either be a 12 gigabyte RAID 1 or a 24 gigabyte non-RAID They then went to a 1.2 gigabyte drive for 30 gigs. I think they went to two gigabyte drives 48 gigs, and then they announced a 5500 in '92. They had sold a thousand of the first system in the first two years, and by '94 which is when Iceberg finally ships--so by the time Iceberg is announced Symmetrix is very strong in the IBM mainframe market, and two years later Symmetrix has 50 percent of the market share by when Iceberg ships. Had Iceberg shipped in '92 who knows, but by '94 Symmetrix was a very established competitor, because they just kept following the disk drive capacity curve doubling every two years and just putting larger disk drives in the same systems to get higher capacity, and with the 5500 they had more drives than just 24. I think they went up to a much larger number for the maximum system.

Moore: You have to give EMC credit for taking advantage of a time where they had an up and coming system and the two incumbents in this space were both floundering. And that door opened, and they jumped in, and they took everything they could, performance guarantees, large caches, and out of that they became, as they are today, the giant of the storage industry. But that was the great gateway for them to get going because the competitors were weak. They didn't have to fight off a lot of competitors. They just had to bring it in. Try it, you'll like it. Here's my guarantee, boom, and it worked.

Gardner: Was Iceberg going to be a large cached machine the way Symmetrix turned out to be?

Moore: What do you mean?

Gardner: Well, one of the selling advantages you described to Symmetrix was the fact that a large cache gave it a very high performance and I was just asking about what was the Iceberg concept, vis-a-via the Symmetrix, was the Iceberg intended to be a large cache. I think at the time Symmetrix probably had four gigs of cache where the biggest 3880-13 you could get might have had a half a gig. It was like eight times larger.

Moore: EMC always had larger caches, so I think to say you could outperform them would be a huge stretch on that. But the StorageTek device was to do space management, and write updates, and snapshot copy, and solve, other than the 10 percent of your data maybe that needed ultra high performance, solve the vast majority of storage problems.

Gardner: So it probably wasn't a large cache machine, at least architecturally...

Moore: Well, I say as large. It was supposed to perform also and do high performance IOs, but probably not to the degree. You're talking a 120 mile an hour car, 150 here, and I think Iceberg was the 120 and 150 would be the Symmetrix. For most apps Iceberg should have performed very well.

Gardner: Now, back in '89 and '90 I don't think RAID 6 was a term of the art. Did you folks invent it? I mean, the Berkeley paper doesn't even describe RAID 0. It just covers RAID 1 through 5. RAID 6 has had a lot of different meanings since then.
[0:40:00, 13:10]
Iceberg, I believe you said, was from the start to be a dual parity RAID 6 type configuration. I was just wondering is that a term or a concept that to your knowledge was invented at StorageTek.

Moore: I don't know if it was invented at StorageTek. HP had a very similar product that they called RAID 6 Plus with log structured files. And I can't think of the name of that disk array right now. Theirs never made it to the market to my knowledge. But I've talked to them several times about their products. They were using the very same concept like that, so I don't know who created log structured files RAID 6 Plus. I do know that at StorageTek RAID 5 was not enough. You have one disk with every parity. They wanted dual rotating parity spread completely around the subsystem, which was a much higher availability and it was a great way to put to bed the reliability concerns of the past ten years of StorageTek. [0:41:05, 14:15]

END OF TAPE 4 / BEGINNING OF TAPE 5

[0:41:19, 14:29]
Gardner: Fred, we were talking about the Iceberg, incredibly innovative combination of redundant arrays, hardware emulation of IBM count key data devices, compression and a log-structured file system. Does that sort of summarize the innovations in Iceberg?

Moore: It was the first occurrence of snapshot copy in our industry on that product.

The two-year delay on Iceberg, I think, as I mentioned, was very costly. Industry analysts, the media people watched the hype on this thing, and all the pre-announcement work done with non-disclosures, and then saw no delivery, and questioned Storage Tech. Could they really deliver? There were all sorts of cartoon-like comments made, that, you know, Iceberg sank the Titanic and it was going to sink Storage Tech. The delays were very costly. It took continued R&D expense to fund that development with no offsetting revenue, and it made the 8380 product line live even longer. The 8380-R product line ultimately became the flagship of that line to replace the market that Iceberg was supposed to address. But the blast to Storage Tech was huge, in terms of credibility in the disk industry.

And finally, in June of 1996, two years after it started shipping, in what I viewed as one of the most incredible marketing deals of all time, Storage Tech and IBM agreed that Storage Tech would OEM the

product to IBM, and IBM would sell one of their competitor's products. This was hard to believe for customers, because for years IBM had been saying everything negative they could about Iceberg, and here, the next morning, they're going to go sell it themselves. This is kind of like IBM never doing in a library and then they announce two of them, and this hurt credibility on both sides. It was also a very huge blow to the Storage Tech sales force, who by and large now could no longer sell a disk product. It took away a lot of revenue opportunity for them, a lot of commission opportunity, and the impact in sales was big, and a lot of the sales reps began to leave the company as a result of that. IBM picked it up. I believe IBM got to almost 30 percent of the market share in the mainframe market with Iceberg on that point.

They did well with it for a while, and then began to phase it out in a few years for their own product. To my knowledge, Iceberg is the only disk array to ever offer full inline disk compression with data. Nobody's done it since. We really have got snapshot copy all over the place now with a variety of implementations, particularly in the non-mainframe world.

But it brought out a lot of firsts, and the ability to store more data than you had physical capacity made it very appealing. The reliability improvements of RAID six Plus Architecture with dual rotating parity were also supposed to take all the focus away from the reliability problems of the past. But what happened in this time frame was that disk drives themselves got much more reliable. So the big fear, that disk failures were the only failure I'd ever experienced, was fading. There were other problems now that caused people to do data recovery. Human error, software error, intrusion, all this kind of stuff that led to data corruption was bigger than a disk failure.
[0:45:00, 23:10]
So the need for this had diminished somewhat in that regard.

But at the end of day, IBM wound up with Iceberg. Storage Tech didn't have any disks to sell. A lot of the sales force changed, and the company began significant changes at that point in time. At the end of 1996, the chairman, Ryle Poppa, resigned from Storage Tech, and that kind of brought to the end of the era of all the innovations that Storage Tech had. They had one left, and it was the integrated virtual tape library with a disk array and a tape library, but an argument there would be IBM did it in parallel, so how innovative it was, not real clear, but it was the last great storage innovation that the company brought out with that.

Gardner: I think in early '95, *Computerworld* rated Iceberg the best of the RAIDs. That would be EMC. I think IBM probably had a RAMAC at that time. I'm sure as a marketing guy, competitive analysis was your thing. Can you summarize the position of Iceberg at that point?

Moore: In the mid-nineties, you really had IBM, EMC, Hitachi, and Storage Tech as your four main players selling mainframe storage like that, and the name of the game was EMC. EMC had gotten tremendous momentum, visibility. They were known as a high-performance storage company, and their products worked, and they backed up their requirement, and Iceberg actually failed to bail out Storage Tech or IBM from falling into second and third place in this regard. Hitachi has always been there as a disk vendor. They've always been about the same market share, always had a solid product, and never made a huge run in the market at all. They seem content with status quo, but even today, 15 years later, that's kind of the same story with this. But by then, the landscape in disk had changed, and it was now EMC's game. They were the new leader in storage, with the false starts by IBM and Storage Tech. And

then, shifting Iceberg from Storage Tech to IBM, EMC emerged from this shakeup in the disk in the DASD industry, if I can still use that term one more time. They emerged as the clear leader.

Gardner: There was a second-generation Iceberg again in the mid-nineties. Anything you can tell us about that?

Moore: You know, I left Storage Tech at the end of 1997, and didn't follow that product much further. I will tell you that Iceberg was used as the core engine to front-end the tape library for the integrated virtual tape, to put all that architecture in front of a shared virtual array that was used with the tape library. The new Iceberg that you referred to had mainframe and non-mainframe attached both to extend that to the non-mainframe world, but the challenge there was that Storage Tech didn't have any sales force that really sold much to non-mainframe markets, so the market share never got to be sizable on that.

Gardner: Did they try to set up a separate sales force, or try to train their salesmen to deal? That's a classic sales management problem in our industry. How'd they try to deal with it?

Moore: Well, you know, then the game was shifting to channels, just to go sell your product, and that was a tougher sell for channel. It hadn't been through all the storage management issues. Channels, more often than not, our price only is our whole game here. "Buy from me, I got a better price." So, in general, the channels were not at all effective in selling that.

Gardner: The STK sales folks couldn't sell against IBM their own product. IBM was just too big, too wide. What was the underlying issue there? As far as tech product, everybody knew Iceberg was Storage Tech.

Moore: Yeah. Well, and part of that deal was, though, that IBM was going to take over and manage the Storage Tech microcode and engineering team, and that all new development enhancements would come at IBM's discretion, that those new enhancements would be in sync with IBM's DFSMS software suite of storage management capabilities, and the product would take on IBM strategies. So IBM had a really good story there that "we're going to do all the enhancements, midlife kickers, new features, and everything will come out at our discretion."
[0:50:00, 23:10]
So what people saw was a platform that they believed IBM would build on and really enhance.

Gardner: I didn't realize that it was a transfer of the design responsibility in addition to an OEM sale.

Moore: The engineering budget that Storage Tech engineers actually became paid for by IBM.

Gardner: But they remained Storage Tech employees.

Moore: They stayed out here 'cause all the-- everything was out-- they stayed in Louisville.

Gardner: As Storage Tech employees?

Moore: You know, I-- and I guess I don't know technically on paper where that came from. I think they were Storage Tech employees, but the R&D funding all came from IBM for the future. That was part of that deal.

Gardner: To give up your future is an unusual decision by management on the Iceberg. Could you explore why IBM, why Storage Tech, drove them to essentially take what had been troubled, but now looks like a successful product, and turn it over to its chief competitor?

Moore: Well, as you probably know, the deal was financial, and it was believed, because IBM had so many more feet on the street, and so much more coverage, that they could get dramatically more market share that it would become financially viable, and the DSS product line, for maybe the first time ever at Storage Tech, would become profitable. Storage Tech did just simply not have enough feet on the street to get the market share for this kind of product. Plus, IBM's future enhancements, to link it in with their operating system and SMS architecture would have greatly improved time to market. So the belief was, this will be a financial plus for the company, a strictly financial deal.

Gardner: But driven by the higher volumes, because of the broader coverage by the IBM sales force.

Moore: Oh, yeah.

Gardner: Did that happen?

Moore: My belief is IBM got it up to around 30 percent of the market share before they moved on to another product line after that. They decided not to keep it alive. It was complex, and ultimately the product agreement, the OEM agreement, ended, and IBM did not renew it. And here Storage Tech, then, had Iceberg back again now, and then this time, the sales reps who wanted to sell it, several had left-- key reps. The system's engineers, who only had tape left, you know, they lost their great things to keep them excited. Iceberg was, you know, one more ultimate joy for an SE [ph?] to go work on, and that was gone. And the infrastructure when the product came back to Storage Tech was too weak to sell it. Product was used as a front end for the tape library, but it died right off.

Gardner: And how big a share do you recall Iceberg got prior to the OEM deal?

Moore: Oh, probably 12, 14 percent of the market.

Gardner: So maybe triple share?

Moore: Yeah. Two-and-a-half to three times, yeah.

Gardner: But you're selling it at OEM price, not at a PCM price.

Moore: Plus, you're getting your engineering paid for.

Gardner: Interesting. Great short-term financial results, and a long-run ticket out of the business.

Moore: Well, a lot of the thought was that IBM would use it as a foundation to build a huge architecture on, and, you know, they didn't. They had simpler paths to go, and the non-mainframe storage market got so much bigger that that's where all the money was, and disk was for non-mainframes, as, you know--the same thing going on here. The mainframe market was under great fire. People thought the mainframe was going to die. Most everybody you talked to wanted to leave the mainframe and get off of it, so the new focus for R&D and disk was moving, to a large degree, to non-mainframe disk storage:
[0:53:00. 27:06] End of …2.VOB
low-cost arrays, JBODS, cheap midrange disk. SATA disks started rolling in. I mean, people had a different economic model for disk once they became reliable enough.

Gardner: Except those folks in the mainframe, where they're still buying EMC and IBM and Hitachi, and some Network Appliance, but Iceberg-type systems are still what are being sold.

Moore: Well, you know, there have been 500 new mainframe accounts since the year 2000. There are about 8,500 mainframe footprints right now. The z/Linux platform's growing at over 25 percent a year. There's growth in the mainframe space. The new zSeries, the z10 [ph?], can contain up to 1500 Intel servers, if you want to consolidate. So the savings are massive here.

And there is still growth in the mainframe. The thing is, it's 8,500 people, roughly, as opposed to, you know, 850,000 customers. So the number of customers, actually, is slightly growing again, and there's growth back in this space. The mainframe was a dinosaur once, but, you know, right now it's a jaguar. It is thin, it's lean, it's 32 square feet on the floor, and I can put 1500 servers in there.

Gardner: And it has probably 60 percent hardware margins. <laughs> It's still, to this day.

Moore: It should. Yeah, it was about a $5 billion business for IBM in 2007. I know '08 was tough with the economic collapse, but that hit every financial number that we know.

Gardner: So you left Storage Tech and have been doing your own business development and consulting. Would you like to tell us a little about Horizon?

Moore: Yeah. I left on January 1, 1998 and came here, where we sit in Boulder, Colorado,
[0:55:00, 1:04]
started Horizon Information Strategies to help customers with business development, and vendors with business development, storage strategies. I write a lot of light papers, publish books, sit on a few boards, and speak all over the computer industry as a speaker. And unlike our tapes here, I speak on the future

and where this is all going, so you'd hear things there like "cloud computing," for example, and a different conversation, but both a consulting and industry analyst role right now is what I do at Horizon.

Gardner: Is there anything you'd like to sum up the Storage Tech era as you saw it-- give us a closing remark?

Moore: Well, the Storage Tech era-- and I spent 21 years there-- was one that was very exciting. It took me through Chapter 11 on how to reorganize a company. It took me through product delays, product failures, how to work with customers, how to build credibility in times of adversity. Probably the greatest learning experience that I ever had was at Storage Tech, to go through all those tough issues, face to face with customers, analysts, investors, and employees, even, and to understand how to make the best out of the situation. So I can look back on those 21 years as tremendous years. Twenty-one years, however, was enough. It was time to go and see if I could stand on my own and earn my own income, and live a life that I could create here without any corporate boundaries. So, wonderful experience, and I never look back about making the decision to start my own business.

Gardner: For the Computer History Museum Oral History Committee, I'd like to thank you very much for the time you've given us. We appreciate it very much. Thank you again.

Moore: Thank you, Tom.

END OF INTERVIEW

