The Disk Drive Story
Chapter 1: IBM's RAMAC

Transcript #2

Transcript of Videotaping on December 3, 2001

Participants:	**Louis D. Stevens (Lou)** Joined IBM – 1949 In Poughkeepsie Joined RAMAC Project - May 23, 1952
	John M. Harker (Jack) Joined IBM & RAMAC Project - May 19, 1952
	Alan F. Shugart (Al) Joined IBM – June 11, 1951 Joined RAMAC Project – 1955
Moderator:	**James N. Porter (Jim)** President, DISK/TREND, Inc.
IBM Project Sponsor:	**Dr. Robert J. T. Morris** Director, IBM Almaden Research Center Vice President, Personal Systems & Storage
IBM Facilitator:	**C. Denis Mee**
IBM Assistant/ Transcribed By:	**Merridy Howell**
IBM Camera/Audio Visual Coordinator:	**Farukh Basrai**
Location:	**IBM Almaden Research Center** **650 Harry Road** **San Jose, CA 95120**

Transcript:

Jim: This is the second installment in the disk drive chronicle. We are trying to recall what was going on in the development of the first disk drive in the world. Today we have, starting to my left: Lou Stevens, Al Shugart, and Jack Harker. Could you please tell us briefly when you came on board and when you got started on the project – Lou?

Lou: I started with IBM in 1949 working in Poughkeepsie. My first three years were spent there, and then I came to San Jose temporarily for the start of the laboratory and to help Rey Johnson recruit. I got a *(permanent)* job and came back out and worked in San Jose starting in May of 1952.

Al: I started IBM on June 11, 1951. I graduated June 10, started June 11; I remember that because I needed the money. I was a field engineer and I transferred to the laboratory at 99 Notre Dame in 1955, having told IBM that if they were looking for someone with field experience, I had seen all the problems fixing machines. I worked at 99 Notre Dame from 1955 until they opened the plant site at Cottle Road. Then I moved down to the plant. I worked on the 305A, 305, mainly the processor, then went to the plant and became product engineering manager for the 305 and left IBM in 1969.

Jack: I joined the laboratory in May of 1952 fresh out of college – Masters Degree from Berkeley.

Jim: In our last discussion we went through a lot of the details on the thinking that generated the need to develop a disk drive at IBM. We're starting this session at the point where a decision had been made to build a production disk drive. Lou could you start us off on a discussion? What created that thinking and how was that decision made?

Lou: First I'd like to back up just a little bit. It wasn't an easy decision to make; to decide upon a production disk drive. We had only just gotten an engineering model operational. Rey Johnson was very adamant about getting somebody to recognize this accomplishment. He did a little bit of backdoor product announcement through PG&E by putting a picture of the disk drive in the *PG&E Progress*. That created a lot of flack. Rey decided the way to get the ball moving was to make several field test models of the 305A processor that went with the early disk drive. There was a movement afoot to get some sort of a field test program, but it was going to require some money from the corporation. A lot of people didn't think it was a very good idea to invest in the field test machines. Fortunately, we had one or two people who worked at high levels who thought it was a good idea. F.J. Wesley worked for Red (L. H.) Lamont, who at that time was Executive Vice President. Somehow Wesley finagled a decision to make fourteen test models of the 305A – one of them went to you (referring to Jim Porter).

Jim: Lou is referring to the fact that I went to work for the company that bought that drive, Crown Zellerbach. I do remember showing off the world's first disk drive to business guests because it was such a good show.

Lou: Zellerbach Paper Company was a very important ingredient in our early work on the processor side of the house. Everything referred back to Zellerbach applications. They were the starting point of the disk drive application.

The decision was finally made to build the fourteen field test machines. Those fourteen field test machines were to be manufactured in San Jose in the building on South Tenth Street by as yet undetermined manufacturing resources (laughter). The influx of a young man from Poughkeepsie, Ernie Freedley, was a key addition to the San Jose manufacturing world, coupled with another later resource – Frank Paul. Frank became the chief manufacturing honcho.

Jim: Somebody collectively decided to build these disk drives. What did you have to do before you could build them? Were you really ready at that point in time with the technology to build these drives or were there problems yet to solve?

Lou: There were problems yet to be solved, but I think for most of the basic problems, we at least had laboratory-level answers. We didn't have production-level answers. We didn't have any experience with production-level stuff. I think there was reasonable assurance that we would find an answer to those problems as we became more aware of what they were. Questions came up like, how were we going to make that many disks and how were we going to do the other things related to mass production.

Jack: By that time the prototype disk drive was operational.

Lou: It had been operational for some time. It was operational at the time the *PG&E Progress* came out, but the system wasn't quite there yet.

Jim: Al, you worked on this. What was involved in getting this thing started?

Al: I was thinking about the difference between the 305A and the 305 – the 305 being the production model. Actually, the disk drive from each was quite similar. It was much farther along in the process. The 305A processing unit, first of all, was a relay-driven logic machine. It was not a stored program machine. It was a control panel program machine. It was really very rudimentary. Somebody else architected the production version of the 305. There wasn't a lot of big change from the 350A drive to the 350. It stayed just about the same. They improved things here and there, but the processor totally changed. We went from relays to tubes; some people even wanted to go to transistors in those days, but we did not do it. The processor was totally different. That's what took the time.

Jim: Were there a number of significant challenges left to build that production drive? Had all of the development been done for all the parts needed to go into it?

Jack: It had been designed. There were a lot of items that needed some refinement. I was responsible for one or two of those. For instance, the pinion that drove the clutch had a life of about thirty hours.

Jim: What were the biggest challenges?

Jack: Getting reliability. None of us had ever designed anything to be manufactured or **used**. We were just inexperienced engineers.

Al: I remember a big challenge was the disk drive because it was an externally pressurized head, air bearing. So you had to have an air pump to pump air into the head. To get a reliable air pump was a big challenge.

Jack: Everybody thought the pump was just fine – every garage has one, but garage-type pump only worked a few hours a day. Something that has to work continuously – it was a new challenge to find that.

Lou: Another question was where to hide the air pump in a machine room. With the RAMAC machine, everybody notices that there is a box associated with the desk. That box contains the air compressor, which is properly baffled to keep its noise inside. It's a continuously operating air compressor that takes about one cubic foot per minute of air to supply the heads. That's a lot of air!

Al: And there were only two heads.

Jack: That's right. There were a lot of challenges on the head/disk interface trying to get the head close enough, yet far enough away, to be safe. The disks really weren't very flat and that was a real challenge. The magnetic recording was a real challenge too. Those rather trivial densities were hard to achieve.

Lou: You used the words 'head disk interface'. That reminds me, we were well into production of several machines, maybe the fiftieth or sixtieth RAMAC and we had a thing called HDI "Head To Disk Interference," as the carriage moved up and down the elevator. When it stopped, from time-to-time it would cock a little bit, just enough to squirt the head into the disk and create a little interference, if you will, a euphemism for **a crash**. The head-to-disk interference problem was not trivial. That required an immediate fix on the design of the carriage to strengthen its backside.

Jack: Again, we were using components that had not been designed for continuous use at a high rate. The potentiometers that we used for the track positioning were never designed to be used in a dynamic way. They were high-precision potentiometers and wire bound, as I recall.

Lou: No, Markite conductive plastic.

Jack: Track-to-track instead of disk-to-disk.

Jim: Al mentioned that the original fourteen test runs were put together with relay electronics, but they eventually changed to vacuum tubes didn't they?

Al: Relay logic was originally used and a control panel would do the control work. Circuit breakers were on the drum. There was a core memory—100 characters.

Lou: Which was hard to come by. At one time early on, John Haanstra's view was that we should synchronize the drum and the disk. We tried that very hard, but failed to get the disk and drum synchronized in such a way that we could make transfers between disk and drum. That was a disappointment when we had to go to some sort of asynchronous buffer. Since transfers on the machine were 100 characters at a time, the 100-character core buffer was adequate, but was not readily available.

Jack: The core planes were hand strung.

Jim: Speaking of typical manufacturing, when you started going into mass production of heads, which had never been done before, surely this must have been a challenge?

Al: I was really working on the disk drive at that point in time, but later on I realized that we weren't working with ferrite heads even, we were working with stack laminations – quite a different head and very very difficult to manufacture with stack laminations.

Jim: IBM was really responsible for manufacturing all of those components: the heads, the disks, and the electronic assemblies. You had to put all of that together before you could assemble the machine didn't you?

Lou: In general, IBM in those days tried very hard to manufacture everything. Perhaps that was an error, but parts of the unit -- the vacuum tubes -- were IBM products; relays were all IBM products.

Jack: Aside from standard parts.

Al: The wire contact relays were all IBM manufactured.

Jim: But you were building a unique product. There were no sources.

Lou: Mostly we used IBM standard parts from the old accounting machines.

Jim: Thinking about it today, to me, a person who didn't work on it, it's a little surprising that you were able to ramp production of these heads and disks, for example, and get production going in an orderly way.

Al: Keep in mind that heads were only two per drive – there weren't a hundred.

Lou: We had a couple of very talented model makers who could produce the needed heads in no time.

Jack: It was a talented group that put it together. Look at the history. A lot of them went on to bigger and better things.

Jim: A new facility was set up on South Tenth Street to do the actual production?

Jack: Because we didn't have room.

Lou: 99 Notre Dame ran out of space.

Jim: As I recall, production of RAMACs eventually reached a grand total of about 1000 machines.

Lou: Something like 1,200 give or take – less than 2,000 for sure. The decision was made in about 1956 to purchase the tract of land for the Cottle Road facility. From there on in, Gav Cullen came from Endicott and the manufacturing organization began to build during the early '50's. The plant was opened in 1958.

Jim: RAMAC stayed in production until the early '60's?

Lou: Pictures of Bldg. 5 exist with the 100,000 sq. ft. open space filled with RAMACs.

Al: IBM wasn't afraid of those sorts of things. When putting the 305 RAMAC into production, at the same time we were putting the first IBM small computer system into production – a small scientific computer system, the 1620. It was a Poughkeepsie product. It was my problem too. We put that into production along with the 305 RAMAC. No big deal.

Lou: It was a big deal! It was a big deal (laughter). You guys were just good, that's all.

Al: I remember the small scientific computer project from Poughkeepsie, the 1620. There was a code name for it because it didn't have an arithmetic unit; it had a table lookup. The code name for the 1620 was the CADET, which stood for **C**an't **A**dd **D**oesn't **E**ven **T**ry (laughter). I always remember that.

Jim: Speaking of logic, what about the RAMAC's original logic? People today think of complex encoding schemes and addressing schemes. It had to be a lot simpler in the case of the original RAMAC 305 didn't it?

Al: I don't think the arithmetic scheme was complicated. Jim Haywood did most of that and it was a fairly straightforward arithmetic system. I didn't have to do that. I was working for Dick Weeks at the time, and was more interested in the general processing of information. We used tracks on a drum. I don't remember it being too complicated.

Lou: I think it is an important point, for the time, computation was not really the most important thing in most clerical work. Computation was not important, but moving data back and forth was. Moving data back and forth was what it could do.

Al: The first two patents assigned to IBM that I applied for and got issued, had to do with moving data. In fact, they were shift register patents. The patent attorney said, "Gee this is a unique thing" – it just came natural. Shift register. Moving data.

Lou: Arithmetic was not important because it was such a small part of the job.

Jim: And the method of addressing each of the files and to be able to find it.

Jack: It was all fixed record length; sequential addressing.

Al: Things were very simple. Things were very very simple.

Jim: Patterned after tabulating cards.

All: Quite a bit.

Al: The coding used on the disk drive was a 7-bit code – 6-bit + 1 binary. Very straightforward and simple.

Jim: I understand that synchronization was quite a challenge to achieve.

Lou: It wasn't the synchronization per se; it was getting the data off and clocking it. That was a real problem. The gentleman who solved that problem was Len Seader. He invented a very clever scheme, which had to have data coming at it all the time, and what was the euphemism, Al?

Al: We called it 'crap in the gap.' You couldn't have a blank space coming into the amplifier. There always had to be some data there. The amplifier was depending upon something being there. So we had to write fake stuff in between the address and the data. Leonard came up with a self-clocking scheme that was really something.

Jack: It was the first. Up to that point all the machines had clock heads. The original proposals for the disk drive were going to be externally clocked from a clock head. I think the first disk drive that was built was externally clocked.

Lou: Clock heads to synchronize an oscillator don't really work out very well. They varied too much.

Jack: There was too much mechanical vibration on the disk and the heads on the clock.

Jim: The method you've just described turned out to be successful then?

Lou and Al: Oh yeah.

Al: From an intellectual property standpoint, we were in great shape because nobody had ever done it before. A clock had done all other clocking of magnetic recording data. So that put us in a great intellectual property situation, as I recall.

Jim: Where did it go from there? You got into production? Lou, to go back to the role of each of the people, as I understand it you had been appointed manager of the project again to get it into production. Is that right?

Lou: Yes.

Jack: The lab split. Rey took the portion of the group that wasn't actually working on the RAMAC.

Lou: The IBM Corporation reorganized in 1956 and established the Data Processing Division. My boss was Ralph Palmer, the director of engineering for the Data Processing Division. Ralph selected me as the new lab manager for the San Jose Laboratory. When that was done, Rey Johnson and his group of people continued with advanced development but they moved from the 99 Notre Dame site over to Julian Street. Rey was still supportive of the RAMAC project and many many people moved back and forth between the two organizations. We were apart at the top, but not apart at the bottom. Rey reported through a different channel than I did, but we were still very much San Jose oriented – RAMAC oriented. When I needed help, I went to Rey.

Jack: That's when I, along with John Haanstra and Hal St. Claire, was commissioned to do a transistorized version of the RAMAC. That was a challenge! Hal left, Don Stevenson joined us, and we were assigned to do a low-cost RAMAC.

Lou: I can't remember at what point it came in, but John Haanstra was doing a big study before he left, which included the Air Force Supply problem and the others that we had looked at.

Jack: It defined the requirements for storage for a fast drum and the ADF, disk drive on a tape ticker drive.

Lou: Which became the 2321 that Al was involved with at one time.

Jim: But in the case of the RAMAC, staying on that for the moment, at some point you moved on didn't you Lou?

Jack: Oh, not for a while.

Jim: You stayed in charge getting it into production?

Lou: Nominally in charge, but I couldn't have done it without the guys that were around at the time. Al was one of the key guys. We had a lot of other key guys who made things happen: Trig Noyce, Don Johnson, and others.

Al: We kept adding improvements to the 305 RAMAC. The disk drive of the 305 RAMAC was a 5 million-character disk drive.

Jim: Let me make the point that characters are not the same as today's bytes.

Al: A double density drive was introduced – a 10 million character disk drive. That was just one of the improvements made to the 305 RAMAC system. New optional features were continuously added; so you could charge more, make it go faster.

Jack: In addition, the disk drive was attached to the 650.

Al: Oh yeah.

Jack: That was a challenge, because they wanted multiple moveable access mechanisms. It was hard enough to keep one of those going by itself, but when you've got three of them.

Jim: Is this when you did more than one head assembly on the drive?

Lou: No, this was three separate independent drives.

Jack: Actuators, independently working at the same time.

Al: It was worse than that because people were looking at dual processing and dual actuators, which meant two access stations on this side, and two access stations on that side, for a total of four access stations; or if possible three and three. People thought access stations were something you just added willy nilly. We didn't do a dual access system, and you know why? Because as a practical matter, once you got one of them going this way and the other one going this way and synchronized, you could roll the disk drive over. (laughter) So we didn't do the dual access. But we tried. We tried dual processing too, but the 650 was a usage for the disk drive and there were other computers that did use the disk drive. In fact, the first Stretch drive, as I recall, in fact I know it did because I did it, had multiple externally pressurized heads didn't have any motion this way -- a head per surface.

Lou: A desperation move.

Al: The Stretch system was for Los Alamos scientific computing. We put a head on each surface. Now, talk about air compressors! We had so many air compressors for all those heads. Fortunately, before IBM had to deliver it, the self-acting air bearing head came out and we didn't have to deliver it.

Jim: Going back to the standard RAMAC again, I have to tell you, as a person on the customer side, the RAMAC put on such a good show through the glass panels on the front. People could see what was going on. That was a great feature. I took dozens of people down to the basement at Crown Zellerbach to show them the RAMAC running. Knowing what I know about the computer industry now, and thinking back on that, every time I took somebody down to show them the RAMAC, it was running. One of the questions that would invariably come up was – what is the mean time between failures?

Lou: You wouldn't ask that. I don't have any real facts, but I would guess the mean time to failure was something like eight hours.

Jim: Remarkable. Because as I tell you, I took dozens of people to see it, and every time it was running.

Lou: There's an old story about people renting apartments in New York. If you only look at them for 5 or 10 minutes the next time the subway goes by they're not there showing apartments.

Jack: RAMACs were mesmerizing to sit and watch. Actually, if you watched carefully, the heads went between the disks, the adjacent disks moved apart, and you could see that.

Lou: Expect it.

Jim: So people did look at the thing on their first visit. They didn't just look the system over, they studied it. It was a marvel.

Lou: It was.

Al: It was fast too for those days. I remember Tom Watson, Jr. coming through the plant in the early days when we had just started production of the 350 disk drive. He asked about the access time, and somebody told him. He checked his watch and said, "It doesn't seem that fast to me." (laughter)

Somebody had to do an analysis of what Mr. Watson was really looking at to show him that it really was fast – about 0.7 seconds.

Jim: It was a lot faster than a clerk fussing around with tabulator cards.

Al: It was really fast!

Jim: Well, I think we've covered today's topic fairly well. I appreciate everybody's time.

Al: Enjoyed it.

Jack: Thank you.