The Disk Drive Story
Chapter 1: IBM's RAMAC

Transcript #3

Transcript of Videotaping on January 11, 2002

Participants:	**Louis D. Stevens (Lou)** Joined RAMAC Project - May 23, 1952
	John M. Harker (Jack) Joined 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 Storage Technology Division** **5600 Cottle Road, Bldg. 50** **San Jose, CA 95193**

Transcript:

Jim: We have an interesting group in this session. To my immediate left is Jack Harker, to his left is Lou Stevens, and at the end next to that very interesting disk drive is Al Shugart. Our objective today is to take up the question of the 1301 disk drive – how it was decided that it should be built, the tasks involved in putting it together, and why it was announced in June 1961 when the project was created in 1957 as the Advanced Disk File or ADF. Lou, I wonder if you could tell us what went into creating that project and how it got started.

Lou: One of the members of the laboratory, John Haanstra, was assigned to study potential future products utilizing the disk drive. The RAMAC had been successful in changing the world of applications from batch to partial on-line. It was clear that direct access memory was going to be an important ingredient of any future product lines. Haanstra put together a small team to study different applications and different technologies. There were several activities going on in the company. The ADF which became the 1301, was one of the three products which was proposed as a result of the study. The study talked about a disk file with floating heads and comb access – two really key ingredients of the 1301 – and a unique coating for the disks and perpendicular recording that we'll talk about a little bit later. The original proposal was about a 200 million-character file.

The second project that was proposed was a Strip File, which became the IBM 2321. The third was a high-speed drum – a million-character file.

Jim: How was the ADF team put together?

Lou: The team was put together basically from the existing members of the Research team. They came to the development laboratory and built the first model. Jack Harker was the manager of the group that built the first model. I believe it was called the Test Bed, wasn't it Jack?

Jack: Yes, the first model built was a result of Jake Hagopian's work. Jake conceptualized the use of gliding heads, or floating heads, or self-acting air bearing to support the magnetic element, so you didn't need a compressor. He then evolved that idea so that, if a compressor was not needed, you could have a head for every disk on each of the disk surfaces, avoiding the need for the xy actuator. That was really a key technical breakthrough in defining the machine.

A simple model of the first one was built and then we set out to engineer a prototype for a production machine. That's what Lou referred to as the Test Bed. That was when the really key work, basically the work of Ken Haughton, Russ Brunner and Bill Gross from Research, was done -- defining what it took to make some of the slider bearings a stable platform for a magnetic head. That was certainly seminal work in the whole industry.

We also built a prototype of a hydraulic actuator -- both a cylinder and piston-adder. It became the model for the actuator that was built and used in the 1301. That file also got committed to a key application. At that time, IBM had a program to build the Stretch machine for the AEC which was to be, in those days, a large high-speed computer and needed a very fast buffer for a lot of data. The idea was to use the configuration of a head per surface, with each head having its own read/write amplifiers, so it could read in parallel and write in parallel. That became a key for later development, because it was a simple takeoff of a commercial product. I think at that time we also committed that the ADF would be the file for the system that IBM contracted to create for American Airlines – the Sabre System.

Lou: I think that's right.

Jack: At this point, when it moved into development, there were a lot of commitments. A lot of basic technology work had been done, but as we learned, there was still a lot that needed to be done.

Jim: There is an interesting point about this product that differentiated it from most of the other drives that followed, in that it was going to use perpendicular recording wasn't it?

Lou: Yes, it was the early model built on the basis of perpendicular recording and a unique disk, which was a multi-layer magnetic surface with a soft magnetic substrate of steel and a hard magnetic surface of Fe_3O_4, black iron oxide. This was a fine answer – worked fine up to a point until you really started to make a lot of disks. Then you found out that the imperfections in the base substrate were such that you could not manufacture large quantities of disks with steel substrates. But the vertical recording worked very well.

Jim: To recap some of the key aspects of the project as it originated; it was going to use flying heads and large disks, following the precedent set by the RAMAC. The project really then got formalized as the ADF in 1957?

Lou: More than formalized; it got formalized and funded. We also got committed to a customer, which really should not have happened until we finished our homework. But we were committed to a customer and we may even have been committed to two customers. Very important.

Jim: Who was running the project at that point?

Lou: It was in my jurisdiction, but I'm not sure who was actually running the project. Was it you Jack?

Jack: No, Trig Noyce was the program manager. As time went on there was the product manager, and then program managers reporting to the product manager. That would have been the very late days of the '50's.

Lou: What happened next?

Jack: You left. Vic Witt came out as lab manager, and I think Dennis Willard was the product manager for a period of time. There were a lot of problems with the product. Al, maybe you could take over and talk about that.

Jim: Put the problems in your lap.

Al: That sounds familiar – 'Al maybe you want to take over.' The ADF project was given to me. I was heading product engineering of the 305 RAMAC systems and all product engineering parts of the IBM world in San Jose. I was asked to take over the ADF program as program manager. At that time it was in deep trouble, but it was fundamentally scientifically sound. At the same time, Jack Harker and Al Hoagland were also assigned to the program, fortunately, or it never would have succeeded. That was about 1960 as I recall, and the product was announced in 1961.

NOTE: Following graduation from the University of Redlands in 1951, Mr. Shugart became an EAM customer engineer in the Santa Monica office. He was transferred to the Riverside office in 1953, and was assigned to Product Development in San Jose as a design engineer in 1955.

Jack: Ralph Golub had the Stretch File.

Al: The Stretch File was one of the perturbations that made life difficult, but there were a lot of perturbations that made life difficult for us – Sabre made it difficult. Those were fun days. We've never really improved the hydraulic actuator of those days either. Marshall Freedman was the guy who made the invention of the hydraulic actuator. It was a piston-adder kind of thing where piston distance all added up to tell you how far the head could go. That had to be proven too. Once you have a hydraulic actuator, you have to have a hydraulic pump. Reliability is very very important. I hear Lou talking about vertical recording and about steam homogenized disks. Those are all real fun things, but none of them ever worked. I don't think vertical recording has ever been implemented successfully in a product. The one fundamental thing that the ADF had was the slider bearing. Jack talked about working on this in Research a little earlier.

In 1959, IBM was gracious enough to publish in their *Journal of Research* all the information about how they developed the slider bearing head, which didn't require external air pressure. You forgot Al Ousterland. Besides providing all the information on how to do the thing, they published all the data on spacer and load. So companies like Bryant, Computer Products, Telex and a few others started implementing slider bearing heads using the *IBM Journal of Research* data.

Jim: Bryant Computer is a company way out in Walled Lake, Michigan. Several of the industry characters who later became very active in the industry were there. They actually did introduce their Series 4000 and it actually came out before the final introduction of the IBM 1301.

Al: I think it did. Interestingly enough, the Bryant drive was more like Jack Harker's Test Bed with a horizontal shaft.

Jim: It was indeed. In fact, it was so large that I can tell you the story accurately, it had to be bolted to the floor or it would walk across the room.

Al: It had 36-inch disks as I recall.

Jim: Centrifugal force is powerful. They did have their problems. This was the first drive, the 1301, which actually implemented a head on each surface.

Al: Yeah, and reliably so.

Lou: And at reasonable production levels.

Jim: Okay, the concept is all there. Gee, it sounds easy to make all this happen. Any problems along the way?

Al: The most serious problem I remember from being program manager: there were 15 systems managers all wanting a model to test on their system so that they could announce it. We couldn't make 15. We had to make them one at a time over a period of time. So I solved that problem by calling all 15 systems managers together in a meeting back in New York and explaining my problem. "You guys all want a disk drive and this is the schedule, now we're all going to vote on who gets them in what sequence." They really chipped in with their thoughts and we ended up with a schedule. It's that kind of problem I solved as program manager – very important to me.

Jim: It's amazing you got all their cooperation.

Al: They didn't have a choice.

Jack: One of the key technical problems you had late in the program – was when you switched from steel disks, which were pretty flat, to the aluminum disk with the oxide coating.

Jim: This was a switch from perpendicular to longitudinal recording.

Jack: Yes.

Al: The steel disk was a fantastic disk. It was thin, it was structurally rigid, the problem was that the magnetic surface was not perfect enough yet and it had too many errors. Mechanically it was a great disk, but we couldn't use it. We had to go back to oxide-coated aluminum. Jack can tell you what the big problem there was.

Jack: The big problem there was – consistently we'd run heads on those disks and they'd crack. Rather discouraging. Again, it was Ken Haughton and Russ Brunner who came up with the concept of measuring vertical acceleration on the disk and putting that into a specification that could be measured. Therefore manufacturing was given control over the waviness of the disk, which the industry termed 'X Double Dot'.

Al: So you could be assured ahead of time that the disk was going to work once installed.

Jack: Yes, and that was probably the last key technical breakthrough.

Al: There was another big technical breakthrough by Dr. Fred Brooks, who I always thought of as the father of the System 360, even though Gene Amdahl gets a lot of credit. Fred Brooks was in my office in San Jose and we had one big problem – the ADF. Every so often there were errors, not very many, but enough of an error so that the disk was useless. It had to be thrown away. I talked Fred Brooks into the concept of spare tracks. I said, "Why don't we have spare tracks so they don't have to take out the disk? If a track is found that you can't write magnetically, identify that track on the disk and route the address to some other disk that has a good track." I tried it out on some other systems managers, and nobody liked the idea. Dr. Brooks, who eventually went to the University of North Carolina, saw the wisdom in that and said, "Boy we could save a lot of money if we don't have to keep taking disks off." He bought into this 'spare track' thing and that really saved the program.

Lou: That's been used since.

Al: Oh yeah, it's still used.

Jack: That was the first time in a sense that you had an index, which is now standard configuration.

Al: Right.

Jack: Standard configuration. The other thing that we did, that we subsequently cursed a lot, was the idea of variable record length and count key data as the format. I don't know if IBM's gotten rid of that yet.

Jim: Those words are still part of the industry's language.

Al: There were a lot of subtle changes made in magnetic recording heads. The earlier recording heads had stack laminations. We changed that to go to ferrite, and even made our own ferrite. That was a big step for us. With ferrite in use, it could get improved – better and better and better by making the ferrite better and the machine better. A lot of little things like that were improvements that came along throughout the development of the ADF.

Jim: How much of a challenge was it, looking at the 1301 here of course, with the head positioning system accessing every track in parallel at up to 25 disks at a time in a set of 25 -- was it ever a challenge to put together an actuator for the system?

Al: Well sure, as I mentioned, this was a hydraulic actuator – it was a piston-adder. The fortunate thing about the hydraulic actuator is it could be tested offline. You found out if it worked or not before you put it into the disk drive. It was a magnificent achievement, not from the disk drive standpoint, but from a positioning standpoint – you could test offline.

Jack: We looked for alternatives but they were basically pretty expensive pieces of hardware. I remember going all over the country and talking to people who were in servo actuators and everything else. They all looked and said, "Gee, you've got a wonderful thing there and we can't think of a way we'd want to change it."

Lou: The head assemblies were really heavy – massive assemblies. If you were going to move it in a hurry, it was going to take some power.

Al: There is another kind of problem with hydraulic positioning that has to do with reliability – because as soon as it's unreliable, it leaks. Let me tell you, that is disaster! The whole hydraulic system has to be ultra reliable – tubes, pipes and everything – you can't have any leaks. I've seen leaks. Customer engineers told me about coming in and seeing a disk drive covered in oil.

Jim: I remember some early problems with corrosion freezing up the machine.

Al: Oh yes, another problem I forgot about had to do with the composition of the oil. We got a supply of contaminated oil, I'm not a chemist, I don't remember contaminated with what, but over a period of time the inside of the actuator became corroded – talk about a big problem. We were already in shipping mode at that point. We first had to figure out what in the oil was causing this, how we could get rid of it all, and how to ship oil from then on to make sure it never happened again. It was a big problem.

Jim: Part way through the program, the demand came up to really make a shipment for the Sabre System, didn't it?

Al: Oh Yeah.

Jim: You weren't really quite ready with final product were you?

Al: I've been in the disk drive business for a long time and I can't really remember ever being ready. (laughter) Point being, you could always use another week -- that's okay – that's just life. I remember that I spent a lot of

time at Briarcliff Manor in New York where the first Sabre System was installed and where we installed the drive to get it running. I spent so much time there that when my wife opened up her first private school in Carmel, we called it Briarcliff Academy, because I spent a lot of time in my early days at Briarcliff Manor. (laughter)

Jim: Did you use that head positioning system on the initial Sabre System?

Al: I don't remember if we changed later, but eventually, Sabre used the 1301 or 1302. I think they started with it.

Jack: No, the different file was the one for Stretch.

Al: The first Stretch File that was committed had an air-fed head.

Jim: So you fell back on the same kind of head positioning used on the RAMAC?

Al: The problem with that was that you ended up with a room full of air compressors. Each head had to be loaded with air pressure and then had to take the air to form the bearing. All these heads required quite a lot of air. It just wasn't a practical kind of thing. I think it was really fortunate that Jack Harker committed the early ADF to Stretch so we didn't have to produce that thing. That would have been a nightmare – an absolute nightmare.

Jim: The recording method on that drive was perpendicular or longitudinal?

Al: Just like the 305.

Jim: Essentially from the 1301 ADF program?

Al: Actually, I'm not even sure.

Jack: ADF frame and actuator.

Al: Okay. Shows how much I remember. I wouldn't even do it. Fortunately, we didn't have to do it.

Lou: To do one was enough.

Al: Yes it was. I remember a room full of air compressors. That's my vague recollection of the first Stretch File and air-fed heads – a big room of air compressors.

Jim: How much of an effort was it to really accomplish trouble-free implementation of flying heads?

Al: I think that was done. I don't think development had to do anything. I think that the people at Research and the Technology people in the early stage of product development did all the work. I don't think implementing it was a big problem. There were several subtle kinds of problems, like Jack talked about, where we had to specify the acceleration of disk surface and that sort of thing. But I think in general, the flying head came off pretty well. I really do.

Jim: It sounds like all together, it was an interesting program that was a little off schedule on final completion though wasn't it?

Lou: Aren't they all? (laughter)

Al: Before we had to deliver the ADF, Jack went off to do the diskpack drive.

Jack: That's another story.

Jim: A storage story.

Al: There were lots of little subtle problems, like corrosion of the bearings and specification of the front surface and the back surface and the curvature – specifying all these things for manufacturing – these were big problems. But compared to what was accomplished with the result – they weren't big problems at all.

Jim: The resultant product was announced as the 1301 with either the set of 25 or 50 disks. It went to 50 tracks per inch compared to 20 tracks per inch previously used on the RAMAC. It went to 520 bits per inch compared to 100 bits per inch also used previously on the RAMAC. A major advance, thirteen or fourteen times, in the recording density on this one project.

Al: Of course, the one thing that permitted us to do this was flying heads, which permitted getting closer to the surface to get higher recording density.

Lou: I must say though that this led to some problems because there's a little actuator – a torsion bar – which loads the head and it leads to something like 300 grams of force if there is a failure. This head then becomes the tool of a milling machine and mills the disk surface quite

effectively. In later days we eliminated that heavy load to a lightly loaded slider bearing.

Jim: A lot of things were simplified over time. The final product announcement was made in June of '61. My understanding was that it actually wasn't shipped until the following year – characteristic of the way announcements were made in those days. As we know today, the result was that it changed the nature of the whole disk drive industry – the first disk drive with flying heads on every surface.

Al: I think it changed the nature of the computer industry.

Lou: I think more than that. It changed the nature of computer business and the applications that you could consider.

Jack: The Sabre System was really the first successful commercial online transaction processing system.

Al: You know what it allowed? Jim you talk about the specifications of how many tracks per inch and how many bits per inch, but that allows the stepping-stone for the next level, and the next level, and the next level. Look at the recording densities that you get now. You have to take the steps one at a time. This was the first major step – the flying head.

Jim: Of course, from the point of view of running a system, the time to get to the data was drastically reduced with access on each head.

Al: Oh yeah, we learned a lot about writing and reading.

Jim: What else can we say about the 1301? It was a great development in the industry and it was an interesting challenge to get it done. Not easy.

Al: Not easy.

Lou: Thank God for Al Shugart.

Al: Not easy.

Jim: Al did get a certain amount of fame in the company for having pulled that one off.

Lou: Well deserved, well deserved.

Al: I'm not complaining about the lack of recognition at all.

Jim: Very good. Thank you, gentlemen. I think we've covered today's topic.