Oral History of Glenn Henry

Interviewed by:
Kevin Krewell

Recorded August 7, 2017
Mountain View, CA

CHM Reference number: X8307.2018

Krewell: Hi, I'm Kevin Krewell. I'm one of the volunteers at the Computer History Museum. And today, I'm going to be interviewing Glenn Henry and his career at IBM, Centaur, and wherever else it takes us. So, to start off, Glenn, can you give us a little background on where you're from, what it was like growing up, and a little bit of where you were born and raised, that general area?

Henry: Yeah, I was born in 1942. And I grew up in the Berkeley, El Cerrito, the East Bay area. I went to school at El Cerrito High School. So, basically it was a feeder school and I went from there to University of California, Berkeley. Our high school was very scientific oriented, very mathematic oriented. I had a good high school career, later not so good. I'll come to that. I skipped a grade and started taking classes at Cal when I was a senior. And I graduated when I was sixteen. So, I was off to Cal at sixteen years old. That was-- the fact that we were taking a class at Cal while we're a senior was sort of unusual. The local paper actually put an article in it about such a thing. Now, it's much more common that high school people can go to college. So, I was interested in computers from sort of day one. Why? I don't know. I remember when I was twelve trying to build a tic-tac-toe computer out of relays. I remember my mother taking me to the library-- engineering library at Cal when I was, I don't know, in high school I guess— to get books on things. I specifically remember reading about a multi-stable-- a bi-stable multi-vibrator, which I couldn't understand. Basically, a flip-flop in today's terminology.

Krewell: But it sounded cool, I guess.

Henry: What's that?

Krewell: It sounded cool.

Henry: Yes. But they taught us nothing about computers in high school. There was no sign of computers. But when I was a junior, we did a field trip up to Cal where we saw the IBM 650, which they had in the engineering department. And I thought that was the coolest thing I'd ever seen. They described it a little bit. And it seemed very interesting.

Krewell: Were either of your parents into science and engineering?

Henry: No, my mother was an elementary school teacher. My father was a machinist, mechanic. But they just stressed education.

Krewell: Okay. And then you sort of gravitated towards computer as you--

Henry: Well, I was interested in computers. I didn't gravitate right then. If I can continue the story, so I started off good, in college at sixteen. But this was Berkeley in the early '60s, and quite bluntly, there was too much distraction for an immature person like me. And I dropped out after three years to work as a helicopter pilot. There's a whole other segment there. I worked as a commercial helicopter pilot and instructor for a year still doing nothing with computers. At Cal, I had a new major every quarter but generally centered around mathematics. Mathematics was my-- a great interest. So, for a variety of reasons, I decided being a helicopter pilot wasn't a good career.

At twenty-one, I left it. And I got a job as a lab technician at Shell Development, which is the research arm of Shell. And about a few months into the job-- this is 1963, by the way. A few months into the job, they said they just got a new computer, an IBM 7040. No one in the department knew how to use it. And so, they wanted me to go down and figure out what could be done with it. So, in 1963, I learned how to program the IBM 7040 and started writing applications for our department, which was the spectroscopy department. We would run spectrums on various devices and analyze it to see what the chemical was. And in about six months, I had my own office. The chemists were coming to me because the computer was such a valuable tool. In fact, I published something like seven technical reports there in my two years at Shell. So, that was my first experience with a computer. And it convinced me that's really what I wanted to do. I'll make a comment, it was hard to learn in those days because there were no books. We had no copiers as we know them today. So, you had to go to the machine room and read the manuals. And that's how I learned to program. My programs were in both Fortran and assembly language which was called MAP for some reason, for the 7040.

Krewell: And so, you were self-taught. There were no courses. There was no--

Henry: No. I'm sure there might have been some somewhere. But no, I just went down to the computer room and hung around, and read the manuals, talked to the operators, that type of thing.

Krewell: And then where did that lead to next?

Henry: That lead to me going back to school. Shell came to me, and their department head said they would promote me to chemist even though I didn't have a degree because I was doing good work. But he also said that I shouldn't accept the promotion. I should go back to school. Along my career, I've had some really wonderful mentors, managers, people that helped me. I think my career is really a function of, obviously, the right time, the right place, and the right people. I was just a lab technician. But I was treated very well, taught a trade, and encouraged to go back to school instead of staying and taking their big raise and becoming a "chemist".

So, I went back to school, not at Cal, but at the Cal-- what's now Cal State University, East Bay because I wanted to-- now that I had seen computers and loved them, I wanted to get finished right away. So, I had

a plan. They had a summer session, which Cal didn't at the time. So, I figured I ought to plan to graduate in a year, which I did. And at Cal State University, which was Cal State, Hayward then, easier for me to say. They had an IBM 1620. And within a very short period of time, they literally gave me the key to the computer room. I was the most experienced programmer there. They had one programming class, which taught you how to write assembly language for it. I ended up later, at graduate school, teaching that. But I just programmed all in assembly language on the 1620. So, when it came time to graduate, the head of the math department-- oh, my degree, by the way, is in mathematics not engineering. And it was real math. In those days, that's what-- there were no computer science classes. Anyway, I got ready to graduate, and the head of the math department came and said if I would stay and get a master's degree, they could get me a research grant to work on the computer. And that sounded like a good idea. So, I was the first at this Cal State, Hayward to get this research grant, which was from a consortium of western colleges headed by UCLA. It's called Western Data Processing Center. So, I got a research grant, two thousand dollars. And I got a master's degree in a year, as well as doing that, as well as working part-time. And it was the part-time job, which I'll mention in a second, that got me to IBM.

Krewell: Okay.

Henry: All right. So, anyway, all I'm doing is computer stuff. Technically my master's degree is in math. It was real math. It was set theory, topology, things like that. But what I was doing was programming night and day, literally. I had the key to the computer room and sort of ran it. At one point during the year, the math department steered a company to me that had come to them and asked for help. It was Fiberboard Corporation. They had a paper plant out in Antioch. And I was hired part time to do two things. One was to teach their engineers programming. So, I got them hooked up with teletypes with a BASIC system. I taught BASIC programming to them. But the other problem they had-- they had exploding milk cartons. Without going into details, they asked me to figure out a way that they could, in real-time manage the moisture in their paper as it was coming down the line. It's a continuous process. Logs go in here. Paper comes out there. And in the middle, the paper goes through a series of rollers. And so, the temperature of the steam, the speed of the paper, the pressure of the rollers, determines the moisture content. If there's too much moisture, later when they try to melt the polyethylene, which was what's coating classic milk cartons, the water turns to steam, and it breaks. So, I very quickly built a model in BASIC, collecting-- because they had tons of data, collecting parameters from the paper machine. And I could predict accurately, given these measurements, exactly which rolls were good and which were bad. Of course, that was meaningless. By the time I did the measurement, the paper was gone.

So, I went to research real-time control computers. And along the way, I discover-- I went to look at IBM. They had a thing called the 1130, but it didn't have analog I/O. And it wasn't suitable. But they whispered to me that they had a new computer, the IBM 1800, which had analog. And it was really a real-time control computer and was coming out any day now. So, very soon after that, I had to drop the project at Fiberboard because it was time to graduate with my master's degree. And I interviewed quite a few places. But one of the places I interviewed was with IBM in San Jose where they were developing the 1800. There was a small group called Small Scientific and Real-Time Systems, something like that. And I

really was interested in that. And it was a small group. I didn't know it at the time but I really like working in small groups. And I had the choice at IBM of even either working on a small group like that or going into the giant 360 programming, where thousands of programmers were grinding away on an operating system for the 360. And I think-- I've always believed that I would have never been seen again. There are things about this small group that helped make me sort of a star at IBM. But anyway, that's how I got to IBM. I turned down higher offers, satellite control, and things like that because I wanted to work on real-time process control computers. It seemed very interesting.

Krewell: Right.

<crew talk>

Henry: Am I rambling too much, Kevin?

Krewell: No.

Henry: Do you want to do short?

Krewell: This is the whole point.

<crew talk>

Henry: Okay.

Krewell: This is the whole point. We want to capture all the-- so that's good.

Henry: All right.

Krewell: All right, so you wanted to be--

<crew talk>

Krewell: All right, so you wanted to be in real-time control systems, thought that was an interesting field to be in. And IBM was just entering that as far as having a scientific computer. And they were based in San Jose. Was that south San Jose?

Henry: Yes, south San Jose.

Krewell: Okay, so you took the offer from IBM. And what happened next?

Henry: The group also had the 1800, which was a real-time control system and the IBM 1130, which was just a small scientific Fortran system. And again, luck, the day I was hired, someone quit or left. And I was placed in a special group which was working on designing the next small computer system, even though I'd just walked in the door. So, I started-- I did a lot of things in the first two years. But one of the things I did was start working on the design of the next-- IBM's next small scientific system. This was going to be called the 11LC. That was the internal codename. It was going to replace the 1130. And so, there was a small group of us, six or seven people. And I was in charge of sort of the guts of the operating system. Another guy was doing BASIC. One guy was doing I/O. And one day-- this is going to be a pattern at IBM. One day, IBM management came to us and said they didn't want to invest the money in a new processor. Why don't we use the processor that they just developed up in Rochester? So, we said sure, maybe, who knows.

So, a guy came down from Rochester, Don Castella, who later worked for me when I was in Rochester. And what they had developed in Rochester was a small business computer. Jumping back a little bit, before this time, IBM's-- probably their biggest business was unit record. This is punch card stuff, key punches, card readers, IBM 407 accounting machines, rooms full of these things because that's how payroll was done. That's how billing was done, accounts receivable. Computers were just now starting to do this. And computers were huge and very expensive at that time. And so, what IBM had developed was a replacement to all that card equipment, a family of computers called the System/3. And it was card oriented. They had a new card, a ninety-six column card, not the big eighty column card you've seen. And it was to move people from wired panel card equipment into stored program. These things had a disk. You had a disk. You had a program. The language they sort of invented to do this was called RPG. A weird language, but it's perfect for doing these commercial applications. And they had developed a computer. And you had an instruction set. And they came and told us about that. And it was the worst instruction set ever invented for scientific. It had no general-purpose registers. It was memory to memory. You had two index registers, no multiply, no divide, on and on. It was just-- it was a joke. But of course, this is a first of, I think, three times IBM pushed me into a bad instruction set.

So, they said, "That's what you're going to use." We said, "All right." And then they said, "By the way, we're transferring the project to Boca Raton." Okay. So, after a year and a half in San Jose, I was sent to Boca Raton. Now, during that period of time, I got two promotions in a year and a half. So, I was doing well. And when I got to Boca Raton, a couple guys had come from San Jose, but mostly there were new hires and stragglers from other groups. But basically, we built the System/3 Model 6 BASIC. Now, downstairs, I have a complete system. This was the first system that I was really involved with working on. And I was the manager of the operating system. And it was the first, in my opinion, the first IBM PC. You

had a keyboard. You had a CRT. You had a printer. You had a removeable disk drive, a fixed disk drive. And it was an interactive BASIC system. Now--

Krewell: It wasn't compiled BASIC. It was interactive BASIC?

Henry: Interactive BASIC, yes.

Krewell: And you were in charge of it?

Henry: Yes. So, it sounds like a personal computer. There was a couple of flaws. It weighed-- when you see it, it weighed about thirteen hundred pounds <laughs>. And it cost, I don't know what, twenty thousand dollars probably. So, it wasn't very commercially successful. But what's interesting, there were two or three people in that group, including me, that ended up doing very well at IBM. So, the project was lost, but the opportunity we had to build it because it was a high-quality system, we did-- our BASIC ran in a 64K virtual memory. But the system RAM was only 8K. So, we were doing interactive BASIC with 64K virtual memory, all-- the virtual memory's all done in software. There's no hardware support at all. This was a-- the hardware was a System/3. They just took the guts of a System/3 and put a keyboard and CRT on it. So, it was a-- I'm very proud of that project.

But of course, commercially, it didn't go very far. Later, the software was reused in another IBM project where they actually emulated the System/3 so they could use our software. That's the thing called IBM 5100. It was a predecessor to the PC. So, we went through System/3 Model 6 BASIC. Then there was the IBM 5100. And finally, you get to the IBM 5150, which, as we all-- which is the PC. Anyway, so I did that for a year and got another promotion. Then I was put in charge of a software group and- a hardware group that was developing-- that was supposed to develop the new GSD architecture-- GSD was the name of our division-- because everyone realized that the System/3 architecture, this simple thing I told you about, wasn't very good architecture. So, we were going to build a subset of the 360. The IBM 360 is a good general-purpose architecture. In fact, it's the reason IBM sort of took over the computer business. It was so good. So, we were going to build a subset of the hardware. And my team was going to build new software because the software for the 360s was batch-oriented commercial stuff. We were going to build an interactive thing. So, we're doing that. We're writing specs. I wrote an APL model of the hardware. And one day, IBM comes and says, "You know, we really don't need a new computer. We have the System/3 thing up in Rochester." And I started giving the list of why that was sort of shitty architecture. And then they said, "Well, you're right. So, why don't we send you to Rochester. And you can improve things up there."

So, I was sent to Rochester after two and half years in Boca Raton with another promotion. In my first four years at IBM, I got four promotions, which was considered a world record. So, off we go to Rochester, Minnesota. I'll put in a segue here.

I met my wife in Boca Raton. We were engaged to be married. So, when IBM said, "We want you to go--" I was going to be a second line manager in charge of a new computer in Rochester. I said, "Well, I have a fiancé over here." They said, "Ah, don't worry. We'll get a good job for her there." Now, she was working for IBM. <Laughs> So, you can imagine what happened. We get to Rochester. I'm a second line manager, rising star. There's everyone there. And I say, "Well, there's my wife." They say, "What wife?" We'd gotten married on the way. So, she ended up getting a lousy-- she got a job. But there wasn't a good job. And here we are in Rochester. She's from Florida. I'm from Northern California. And we'd never seen anything like the weather there. However, that was my longest stay at any IBM location, which we'll get to. Shall I continue?

Krewell: Sure, that's ironic, though, right? The worst weather of all the locations you were at, and now you're stuck-- you were there for the longest.

Henry: Right, it was the project I'm getting to, which is sort of--

Krewell: Was there any reason why they were obsessed with keeping that instruction set alive when they had a perfectly good 360 instruction set that they could subset?

Henry: Well, no.

<Laughter>

Henry: A lot of politics, and of course-- the System/3 was doing very well. It sold very, very well. It was the lowest cost. They had many models. But it was the lowest cost IBM system. And it did a good job. If you wanted to do billing, inventory control, sales analysis, payroll, you name it, in a small company, it was perfect. So, we're in Rochester. And I was the second line manager of this system software, languages, utilities, and operating system, for a new computer. And this computer was a commercial computer. It was designed to be sort of a one-person version of the System/3. The System/3 was a fairly big computer. And it was card oriented, even though it had a disk. This was a computer that was going to be a desktop thing. But it was going to be commercial. It had a keyboard, a very small CRT, a printer, removable disk drive. It had, you know, it was going to be very low cost. And the--

Krewell: No terminal supports?

Henry: No. This was before terminal supports. They had Bisync, which was a sort of batch communication line. But it would run RPG, again, for small businesses. It-- to keep it low cost, they weren't using the System/3 CPU. They had a-- the engineering group had developed a very primitive micro-processor, right, micro program. And it would do all the I/O control. It would emulate a System/3

and do I/O control. So, you'd emulate the System/3 instruction set and run the operating system of the System/3s. It had to be modified. That's what my group was doing. I looked at that and said this is crazy. The System/3 is really slow. The System/3 emulation was very slow. You've got this thing. It's a lot faster. It talks directly to the disk. So, I took on the project myself, personally, of rewriting the guts of the disk operating system in the microcode. Fortunately, the engineering group was amenable to this. And so, this was a later theme of mine at IBM - moving stuff from high levels into the hardware. When I say microcode, it was just a sixteen-bit vertical instruction set. Nothing hard to write, very simple. It had I think-- it had sixteen registers. But--

Krewell: Was it a monolithic chip--

Henry: No, it was a whole board of chips. By moving the disk operating system into microcode, we made it a lot faster than if it had to emulate-- run the disk operating portion of it in System/3. Most of the software was still in System/3. But the key performance element was going to the disk and back. I remember watching the 1972 Olympics, so that's where we are, sitting in my room coding this. My son was born in '73. So, I remember it was just about the time I was done working on this. So, this is the '72/'73 period.

That got me my first award at IBM. IBM has a sort of awards system. And you get to go to a recognition thing once a year. You get to bring your wife, which is nice. IBM pays for a wife to come. And that was the first time I got an award. That project, that computer became the IBM System/32, which, before the PC, was the largest selling IBM computer, a very good small business, targeting very inexperienced users, costing maybe a thousand dollars a month, which is dirt cheap in those days. I have one downstairs [in the Centaur offices], a complete one. So, near the end of that project-- I guess '73/'74, management came to me and said, "We want to build a new computer in the System/3 family that's much more powerful," etc. The original System/3 had only a sixteen-bit address, so sixty-four KB of memory, and no multiply, no divide, very limited. They had-- at the same time I was doing the System/32, they had done another version, the System/3 Model 15, which had extra memory registers. So, it could address more memory. But it still was a very limited instruction set. And the software was very limited.

Krewell: There was no virtual memory, just--?

Henry: No, very simple. So, I was given the job to get a group together to start designing what we wanted the architecture and the software to be in a brand new computer. So, we were sort of given for a while a clean sheet of paper. So, we went wild. Fred Brooks talks about the second system effect. You know you design a system. And then your second system, you overdo it. You try to fix all the things you didn't do right in the first system. You weren't smart enough until your third system. This was my second system. I remember, we-- there was an ACM article or something about the ten great steps forward in computing. I remember we sat there and read that. And we ended up doing seven or eight of those things. I'll describe

what it is. This is the most important IBM project. This is what got me to be an IBM fellow. There's another one after this, etc.

But basically, we decided that we would build something totally different than System/3, other than it could run the same languages. We decided to not only do virtual memory, to do a single-level store that contains everything, files, you name it, are all in the virtual address space. There's no distinction. So, where today, you read it from a file-- you read the program that's in a file exe. And you load it into your virtual memory area, which is another area of memory. We executed it where it was. Everything is in virtual memory. To do that, we had to have a huge address space. And so, the hardware actually had a forty-eight bit address. This is in the mid-1970s. I mean it's amazing. Beyond that, we said, we're going to be highly secure. We spent a lot of time. At about this time, the government was issuing a-- I forget the color of the books, the orange books? You now, books on secure operating systems, what they wanted, FIPS requirement. So, we said well, the most critical thing to do is protect pointers, addresses. So, we said we'll do capability addressing. That is the hardware had an extra bit at every sixty-four bits. Or was it thirty-two? I don't even remember now. But that extra bit could be accessed only by the operating system. And that indicated it was a pointer. So, you couldn't add to a pointer. You could call a system function to manipulate a pointer, right?

Krewell: Right.

Henry: Then we said there's this new thing we read about being done in the San Jose Research Lab, which is a relational database. IBM databases is, of course, was a big deal at IBM. The databases were totally different then. And this thing sounded very cool. A relational database - you know, a table format. So, we said we should have one of those. And then I said-- and this was extremely controversial, this caused a seven-year argument. But it turns out to be one of the-- it still is important to IBM. I'll get to that. I said there will be a high-level machine interface. We will not expose the hardware instruction set. So, we built a high-level machine interface. And it was really high-level. It was object oriented. That was another thing. I'm just going over it, protected pointers, a forty-eight bit single level store thing. And the hardware's down here. We had a considerable amount of "microcode", but it was just written-- you know, it was written in C-- it wasn't in C. It was PL/1, our system language, up to here. And my idea was that, of course, that the hardware's going to change every year. Every two years, could do more. So, why change your software, right?

Krewell: <Inaudible> abstraction layer.

Henry: Yes. But that was very controversial. People didn't want it because it was going to slow us down. It did slow us down. It slowed down the first machine. But that interface survived. It's still there at IBM. I'm going to tell you the origin of it. So, the concept worked. We made that interface really high-level. We divided what you think is an operating system, and the stuff above, called the CPF, control program facility, and the stuff below which we called VRM, virtual resource machine, or something like that. And I

said well, let's just go do this. Now, at the same time we were doing this, IBM had decided to replace its entire 360 line. And they had a giant-- thousands of people working on architecting a replacement. And amazingly, they had the same sort of high-level concepts. But they had a much more complicated thing, right? So, in one sense, if you read some of the articles, they'll say that what we did was build a subset of that system. But we really came to it independently. Its name was future systems. FS. It was a giant effort replace the 360s, all of them, right, hardware, software, thousands of people working on it. And we're sitting there in the cornfields of Rochester doing this thing.

Krewell: And that's Rochester, Minnesota.

Henry: That's right, Rochester, Minnesota. Yes, not New York. So, this was controversial for lots of reasons, and all along the way we had audit after audit. IBM has a strong focus on audits where they bring in experts from all of the fields. They come. They would come sit in my conference room for two weeks. I'd present charts. They would argue. They would fight. They'd make a report. Alright?

Krewell: Mm-hm.

Henry: And so, anyway I spent seven years doing that. I managed all of the software, which at the end was about 600 people which is a decent size group. So, I'd come up with the original concepts, or I and this team, small group. They became my managers. We grew and we grew and we built this software, was like an control program, the VRM. Also had languages, COBOL RPG, and a pile of utilities - sorts and stuff like that. It was announced in the fall of '78, and then we missed the ship date, which is a sin in IBM of great proportions, because there's a long lead time to these kinds of systems.

Krewell: Yeah.

Henry: Especially something totally new. This is embarrassing to admit but I'll tell you, one time I was talking to my manager, Brian Utley, who was a really good guy and big IBM manager. I was in charge of all the software. There was a hardware manager and he was sort of the system manager, and we were talking about compensation. He said "You know, if you're going to get more money you need to get this thing announced." So, I went out to get it announced and the way you do that, there's a separate group at IBM, a system assurance group that has to sign off on everything. IBM has a process manual this thick. First you do all this. Then you have reviews. You have specs, people sign off, and we just ran over to system assurance, right?

Krewell: Oh.

Henry: Every objection they had I had 13 proofs why that was a stupid objection, right, and by God, we got it announced. This is a big deal at IBM to announce a new system, because this thing's nothing like the System 3. The software's totally different. The System 38 announced the summer of '78. We immediately then realized that we weren't going to ship it in time I can't even remember now, but we ended up having to slip the delivery date a year, and that was a big deal. My boss was called up to Armonk. I was called up to Armonk, and I still remember one of these things, sitting in a room and out there all IBM vice presidents, lab directors, vice presidents, and my boss, Brian Utley, is there and the group above there is questioning me "What do you think? You think you can do that? When can it be," etc., right, and I'm trying to answer the truth and Brian is speaking up. I remember Chuck Branscomb who actually has an oral history..

Krewell: Yes.

Henry: Great, great, greatest IBM manager I know.

Krewell: Yeah.

Henry: Chuck Branscomb spoke up and said "Brian, keep quiet. Let Glenn answer."

<laughter>

Henry: "Tell the truth, Glenn." So, I told them because Brian was trying to say it was a three-month slip, maybe a six month slip. I said "No, it's a year." At this point if we're going to slip we've got to do it right. Alright? We're going to take a year and we're going to do it right, and Chuck said "Okay, then that's what we're going to do," and they all went without me into the corporate management committee where they basically beat Brian up. Brian wasn't fired but he got a-- he was sent to the penalty box for two years. Later he came back and had a big IBM job because he's very capable.

Krewell: Yeah.

Henry: And every day the lab manager at that time would come by my office and he'd kind of say "You know, Glenn, I think maybe we ought to fire you today. Nah, I guess you're too valuable. We won't." <laughs> This went on for a year.

Krewell: Just to torture you for a year.

Henry: Yeah. But anyway, that system, the System 38 was shipped. It was underpowered. The hardware was underpowered. We had to make-- we made it run at 128K, which was big for the day but was small for the software. Later of course they added more hardware, more memory. They added more hardware later in the '80s. It was announced as a new system - was announced, this AS400, same architecture. The S400 later in the 2000s became this Series 1, I guess. Later that became the P System 1 or the P Series 1, something like that. That architecture if you look it up is still around. It's the same characteristics. The instruction sets grown. The address has grown from 48 bits to 64 bits now at the high-level machine interface, but the original interfaces are still there.

Krewell: So it was a scalable design designed to-- it lasted a long time.

Henry: Yes and it's still going. The AS400 was very successful and that, I mean it's hard. In retrospect, it's easy to say I was right at the time. The concept of not writing to the iron was just amazing, right? People would do crazy things to save memory while we were doing the opposite. We were <laughs> eating memory and we keep saying "Okay. Okay. Maybe it'd be a little underpowered in day one but the software won't have to change. It's not the hardware. It's the software, stupid," except you don't say stupid at IBM but that was the thing. Later -- anyway, we're up to 1980. The thing ships and my wife says she's leaving Rochester. I could come or not come.

<laughter>

Henry: But then remember the nine years total now.

Krewell: Yeah.

Henry: We moved in '71 and IBM told me they would move me to wherever I wanted to go inside of IBM and make a job for me or find a job. So, I moved to Austin and they made a job for me, and I will stop here except one more thing. Two years later, I got a very significant what's called a corporate award which is big money, several tens of thousands of dollars which was big in the 80s..

Krewell: Yeah.

Henry: ...from IBM for the single level store and high level machine concept, which were really-- I was the driving force there and that of course became what was the significant factor. There's one more system I'm going to tell you about at IBM that also I think factored into getting me to be an IBM fellow, but that System 38 was certainly a significant step in there.

Krewell: If you want to take a break let us know.

Henry: What's that?

Krewell: If you want to take a break let us know.

Henry: No, I'm fine. So we're up to 1980. I moved to Austin.

Krewell: And if you moved to Boca Raton maybe you would've been in the 5150 group.

Henry: Yes, but I didn't. I moved to Austin because I'd been in Boca. Austin looked like a good place to live, and it was, but there's no job here. So they said "Why don't you be our Ad Tech Manager?" Ad Tech is Advanced Technology. It's always a group in the lab. Working at IBM's sort of nice. In the laboratories, they have product development groups, always have sort of their own research or advanced technology group. You do an Ad Tech group, and by the way, here's a group over here that's building something we don't know what to do with, so you take it," and that thing that they were building that they didn't know what to do with was a RISC processor, a small 32-bit RISC processor called ROMP for Research Office Multiprocessor or something like that. So they were just building a RISC processor because IBM research had done work on RISC in the late '70s, a thing called the 801, but they'd never shipped any RISC. So, this was just a bunch of good technical guys doing this thing. So, I gathered up that group and then I said "Well, what are we going to do with it?" So, I said "What we're going to do is build a engineering and scientific workstation." There were no such things in 1981. By 1984 several companies had built one, a company called Apollo and Sun. But basically, the concept was again, a workstation. It was sort of a-- they used the term 3M was the best word. A million pixels in an all points addressable display. That was fairly new. One MIP of computing power and one meg of memory. Alright? 3Ms. That was your engineering and scientific workstation. That's what Sun and others introduced in the 1980's. We were going to do one. We were going to be first.

Krewell: Silicon Graphics also?

Henry: Yes, they were later.

Krewell: Yeah?

Henry: Yes. In the late 80s there were several of these things out there. The original Suns were a good example. Yeah, Silicon Graphics, a couple other companies have gone out of business, but we were going to be the first. So, I had a processor, but I had no operating system. I'd just spent seven years

managing an operating system. <laughs> So I said, "We're not going to build an operating system. We're going to find one." Well nothing that IBM had was any good. In the meantime, this FS project they had to replace the 360s that died of its own inertia, thousands of people, and they're back to building 370s now. So, I said "Let's go and let's pick up this thing I read about called Unix." Unix is just very-- few people at IBM had heard of it in 1980, '81, alright? It'd really just become sort of visible on like I think '78 or so.

Krewell: Yeah. Data General, is that the company you're thinking of?

Henry: What's that?

Krewell: Data General.

Henry: I know Data General. I don't think they have one.

Krewell: No.

Henry: Oh.

Krewell: The scientific workstations.

Henry: I don't think so. I build more of a mini computer. It wasn't a workstation.

Krewell: Yeah. Right.

Henry: Yeah, they were mini computers and DEC was doing mini computers.

Krewell: Yeah.

Henry: So I said "Let's go get Unix," and then I said "Well we need a partner for applications." So, we ended up going to Carnegie Mellon University and they were our partners, and here's an IBM story for you. So, I got this thing. There's a RISC processor. I don't have this thing. I've got a proposal, charts for a RISC processor <laughs> a Unix operating system which nobody but me had heard of at IBM, and we'll build an engineering workstation, and remember IBM was very, very commercial and business oriented and not engineering oriented, and in particularly Austin at this time was the home of office products division and they were building a thing called the Displaywriter but basically a word processor.

They had another word processor. They were building word processors, right, just to stand alone in networks but specialty systems and I want to build an engineering workstation system. And so, yeah, I'm the manager of Ad Tech but to do this was going to take more money. Alright? I've got to build hardware. So, I went to my boss and I said "Well I need some more money for this great idea." He said "Well I don't know. Why don't you go see my boss?" I literally went from boss to boss until I'd reached John Akers who was one step below the IBM CEO and a year later was promoted to IBM CEO, and I remember now. This is just one of those memories you have. I don't remember much of it. I remember being up in White Plains somewhere or maybe Armonk presenting to John Akers this idea I have and the need for some more money, and I'm making foils and in the back, John's sitting there and in the back of the room there are all the bosses between me and here, about four of them who'd all turned me down. Alright? <laughs> And I get to the end and say, you know, the final chart of course. I say "Here's what I'm going to do. I'm going to build this. It's going to cost this. I'm going to do in this time," alright? And an interesting management technique, Mr. Akers reaches over, grabs that foil, folds it up, and puts it in his pocket saying "You go do that and I'm going to remember what you told me." I'm sure he forgot about it but he became IBM CEO, and what's interesting is I never got any change to my official budget. I just went around after that meeting and said "Well I've been to Mr. Akers and he said build it. So <laughs> get me two of these and one of these," and I collected 50, 60 guys in my building. We had a deal with a commercial vendor who I can't remember now to do Unix, that we were going to modify Unix and they would provide a base Unix, and the hardware guys were building the processor. I'm going to build frames and covers. It was going to be an engineering and scientific workstation, and along the way the rest of-- so I got 50, 60 guys. There's 400 people over here in Austin building a follow on to-- oh, it was the Display Writer. That was the name I think. That's a standalone word processing system.

Krewell: Right.

Henry: Four-hundred people building a follow on. We would roll our eyes and say "Glad we're not them." One day I got called in my boss's office and said "We're really going to help you, Glenn. We think your idea is good. We're going to kill that project over there and we're going to give you these 400 people." It's the kiss of death. "And then, by the way, I know you're building engineering and scientific but why don't you make it a commercial system, too?" <laughs> So 400 people that I don't need nor want were probably pissed <laughs> and why don't you build-- make it a commercial system, too? So, we spent an extra year and a half from the schedule we would have made with the engineering and scientific thing putting in more features so it could be a commercial system, too.

Krewell: Mm-hm.

Henry: And by the time we'd done that we were a year late to the marketplace. This thing became the IBM RISC Technology PC, or RT PC. You can look it up.

Krewell: Yeah.

Henry: It shipped, but again, it's a forerunner of something else. Okay? This was the first IBM system to ship with a RISC processor and with IBM Unix, which we called AIX. AIX is now a huge thing. That's what the IBM systems run on, mainly the mainframes because we'd taken Unix and Unix was pretty rough in those days. <laughs> We'd put in all these user interface things and you'd tell it to do this, that, and the other thing. So at the start, I was just managing the software. By the end of this I was managing the software and the hardware, and again that system was successful in itself but it begat two things. It begat AIX which became important and it begat-- it was the first RISC system and it was followed by a project here in Austin to do their power architecture, which is our [IBM's] current RISC thing.

Krewell: Mm-hm.

Henry: And I was planning-- I was assuming I would get to manage that, but one day they called me in and said "We're making you an IBM fellow and we're going to bring in Andy Heller to manage that other thing," which was fine because I wanted to do IBM fellow work, which I'll come to in a second. So, I was the manager of four systems at IBM, the System 3 Model 6 Basic no one ever heard of, <laughs> System 32 software which was IBM's best seller before the PC, the System 38 architecture and software, which was a big deal. And the RT PC which was again from one of my ideas, which wasn't per se that big a seller but it led to two other things. In fact, late in that project in '84, '85 we created the IBM AIX Project Office because it was clear that IBM was going to be good software, right?

Krewell: Yeah.

Henry: And my group was building-- when I left it to be a fellow, my group was working on very small kernels, you know, micro kernel-like things to run underneath Unix, what have you. Sort of again a forerunner of the future. So that's my IBM experience, 21-- oh. I haven't finished my IBM experience. I just was made an IBM fellow, 18 years. What am I doing here? Okay. So, my two fellow projects. The first was an audio/video visual computer. Now let me explain something.

Krewell: Could you explain the program, the fellow program?

Henry: Okay. Yeah. IBM fellow is a great deal. It's a lifetime appointment. You can work on whatever you want to work on. You can get a budget. The budget has to be negotiated but it's not just you. Most fellows have a small group of people working on it, and that's it. You work on whatever you want to work on. Lifetime appointment, right? And it has perks. You get to go every year. In the meantime, I'd been to three of the IBM yearly recognition things. You get to go to the recognition thing every year. My wife loved to go to those things. They were always held in the good places and stuff like that. So, it is a good deal, and in those days it was much more picky than it is today. Today they make five or six fellows a year. In my day it was two or three or four. In the year I was made fellow, one of the other person who's now by the way a Professor Emeritus at UT, was the research chemist that developed a new thing. The

other person, I can't think what their-- I should explain. IBM fellows almost totally in my day came from research. They were real people. The year after I was made a fellow, two of the IBM fellows won the Nobel Prize, that kind of thing. They were real research people. I'd never set foot in research. In fact, I didn't get along with research. A whole story there. I just developed products but my products had very innovative things in them that became important, but it was very rare for a product developer. I wasn't the only one but that was rare. Most IBM fellows were research people and that was sort of the career path if you're a doctor in research. But anyway, when I was doing the RT PC I remember I'd have to go to management reviews. I'd stand up and say "Well here's what I got and here's how fast it is and here's what it costs," and right after that Don Estridge would stand up and say "Well I've got a new version of the PC and here's how much it costs and here's how fast it was." It was like half the price but we were 15% faster or something like that. So, I became convinced that PCs were the way to go. So, when I became a fellow, I moved myself into the PC group. No other IBM fellow would get close to PCs, and the whole management chain above me. At that time, as a fellow, I'm fairly high level. My manager was in New York and he didn't know a PC from anything. IBM thought PCs would be interesting as terminals to their mainframe business. It's crazy, but I saw what I thought was the PC opportunity. So, the first fellow project I did was what we called the Audio/Visual PC and the idea was that a PC-- this is 1985, '86 which PCs were pretty primitive. The PC ought to be able to play video, ought to be able to play sounds, ought to be able to capture sounds, do all those things that we now take for granted. They didn't do them in 1986. So, we built-- I had a small group of people. We built the prototype. There was a PC here and hardware here <laughs> that did sound and hardware here that did video and we took it to Armonk and I showed it to John Akers who's the CEO.

Krewell: Right.

Henry: And he looked at it and went "Well, that's nice." <laughs> They were worried about other things. I couldn't get any product group interested in it, but so we took it to a point where, and again it was the same thing. Said "Yes, look it, there's this hunk of hardware here doing sound. Next year that hunk of hardware's one board. The year after that it's one chip. Can't you guys see?"

Krewell: Yeah.

Henry: So no one wanted to build an-- so I wanted someone to build, the PC group to build an audio/visual computer. They didn't want to build it. So, then I had the idea that leads directly to my life today, 30 some odd years later. I said "It's clear again that the PC is going to take over the world. And that means, ugly as it may be, that the Intel processor architecture's going to be very important." Remember, I'd just done a RISC processor. I used to have to give charts on how "Look at that 8088 instruction set. It makes you sick, doesn't it? Look at my clean instructions," and I'd gotten general purpose registers. So, I knew, but I decided that processor architecture was going to be key because that's what software's been running on and software was going to be key. So, I decided that what IBM needed to do was take over that marketplace. At the time, 1985, '86 IBM had probably the best

processor design group in the world. They had great technology out of Burlington. They had really smart people. They were building this giant RISC processor. They'd built processors for years, and at the time people don't remember now. Intel was in trouble. IBM bought 20% of their stock in 1985 to prop them up.

Krewell: Right.

Henry: So I said <laughs> why don't we just build a better 486? We didn't know it was a 486 but build a better..

Krewell: Yeah.

Henry: ...future X86? And then I said that the way to do it is take the RISC, that Power PC, power RISC thing you're building and modify it. This is going to sound familiar, modify it so it can also execute X86, and that's what we were going to do, because across the hall and actually in another building, was the main product group, a couple hundred people designing the power processor.

Krewell: Right.

Henry: By this time I'd gathered up 50 people again and we were just-- didn't have to build a whole processor. We just had to design new instructions and try out the mechanisms, things like that so that this thing could run both, and my idea was an intercept strategy. It could run the old, tired PC crap but it could also run the new, spiffy RISC stuff which I assumed IBM would develop software for, etcetera, right? Well IBM really hated that idea. Their audits came again. For a lot of reasons IBM was a "let's do it themselves". They're going to build OS2. They're going to build a RISC processor. They're going to build a-- what is the name of their..

Krewell: Microchannel.

Henry: Microchannel. Everything's proprietary. Everything's better technically. Everything-- but that better technically didn't matter on any of those things. So, comments in the OS2 thing. So as an IBM fellow I was called to give talks and things. So, I was called to Boca Raton to give a talk to the technical community on the future of processors. So, I had charts showing processors up and to the right, you know? And somewhere in that talk I said, because I just wing it, as obvious from this, I said something like "And you've got to do virtual memory if you're-- anyone who builds software in 1986 without supporting virtual memory is crazy." A gasp from the audience because that's all the OS2 developers and the original OS2 was done to run only on the 286 which was non-virtual memory. Even though Intel's shipping the 386 and Compaq's got one.

Krewell: Yeah.

Henry: In fact, I was called again sort of as a technical expert to go to an IBM sort of investigation on how come Compaq's got a 386 system and the PC group does not?

Krewell: Mm-hm.

Henry: Compaq shipped the 386 first and the answer quite clearly was the IBM mentality that sold all these 286s, so they had to support the people. They couldn't obsolete those 286s, and they didn't realize it's a new world and you'd better be obsoleting things, right?

Krewell: Yeah. The 286 was kind of-- its management of the virtual space was pretty braindead..

Henry: It's segment oriented.

Krewell: Yeah.

Henry: Segment registers point to different a memory segment. It's terrible.

Krewell: Yeah. Oh, by the way, just before we go ahead, that mythical Power PC, like to emulate X86, was that the 815 project people?

Henry: That, I'm coming to that.

Krewell: Oh.

Henry: Yes. Uh-huh.

Krewell: Sure.

Henry: Okay. So, I'm now up to a hundred people building this, and I should mention that the key executive in our area whose name, Jack Keeler, a very big IBM executive, he'd gone to school with Andy Grove and IBM had invested in them. So, I was fighting against the tide. So, this went on until '87 and along the way there was a key individual. So, I felt I needed some very specific expertise and it turned

out that IBM was building a 386 in Burlington under license to Intel but they had designed it themselves looking at all the Intel IP, right?

Krewell: Blue Lightening I think it was called.

Henry: Yes, Blue Lightening, exactly. One in my museum. I went up to Burlington, and I hired the design leader of that group, the logic designer. I remember thinking at the time "This guy's either a genius or <laughs> he's pulled the wool over me." That guy is here today at Centaur. He's our lead architect, one of the best engineers I've ever seen, a person by the name Terry Parks. Anyway, I'd built a group of good people. People like to work for an IBM fellow project but it was just getting tiresome and tiresome fighting against the tide. And then one day, we're now in early '88. My wife reports that they think there's a drug dealer moving into neighborhood. <laughs> Well we live in a good neighborhood and the most expensive house, which is only one away. There's us, another house, and the most expensive house in the neighborhood. It's a small neighborhood. Some 23-year-old kid has bought that house. They don't know anything about it, but I mean what would a 23 year old do? <laughs> Where would he get the money unless he's a drug dealer? Well of course that was Michael Dell, and he had just founded in '85 Dell Computer, or '86. I'm not clear and he was just in process in early '88 of going public, but anyway, he and I ran into each other in walks. We would walk the neighborhood, right, and he would talk about his philosophy of doing things and I would talk about what I was doing. And remember, I'm interested in PCs and he had an awfully good story about <laughs> PCs. So, one day in June or July, I forget the-- I just get so disgusted with IBM I say I'm going to quit, and I was like only the second, at that time, second fellow ever to quit. I'm going to quit and go work for Michael Dell.

<laughter>

Henry: I tell the story my wife faints. Well she didn't faint but..

<laughter>

Henry: ...I mean leaving the IBM fellow job..

Krewell: Yes.

Henry: ...<laughs> and going to work for a 23-year-old, but that's what I said I wanted to do. IBM made me go up to Armonk to talk to the executives. Why are you doing this crazy thing, right, and blah, blah, blah, and put pressure on me not to leave. That was interesting. They were very upset seemingly that I was going to leave but they didn't like my project, my PC oriented thing, my pseudo X86. But anyway, in July of 1988 I quit IBM leaving a great job and go to work for Michael Dell and my job was to be the first

manager of the R&D group, the product and development group which I'll say a little bit more about. And my project, I had people working on it. My project became the 815.

Krewell: Oh, okay.

Henry: And it lingered for a while and then died.

Krewell: Okay.

Henry: But I had people working on the project. In fact, one of the persons working on my project in addition to Terry, there's another person here, a first-year person that worked in my fellow group.

Krewell: Yeah.

Henry: And several of the other people here worked on the 815 and later worked on the 815 but it went basically nowhere.

Krewell: Right.

Henry: So..

Krewell: You're now at Dell.

Henry: Now at Dell. So I go in..

Krewell: Do you try to hold on to <inaudible> PC?

Henry: <laughs> So I find out what the R&D department is. There's six or seven guys working directly for Michael. They have no department, no managers. It's just guys, engineers he had. And at the same day I started, there was another, a new college hire from UT. He had just hired a guy by the name of Darius Gaskins, who's also one of the founders. So, Terry, me and Darius-- anyway, I'd go in and there's--

Krewell: Now there's--

Henry: Centaur.

Krewell: Centaur.

Henry: I believe they had. Go in there and there's 6 or 7 guys and Michael says, "Get them organized. Make a group and I want to see us beating Compaq." I said, "How we going to beat Compaq with 6 or 7 guys <laughs>?" And Michael said, "You're going to be lower cost than them, and you're going to run faster than they do. You're going to turn out new innovations quicker." And that, of course, always stuck with me and that's one of the themes of starting the company I'm in now, Centaur. So anyway, I gathered people. I hired people. We grew first 3 years at Dell. Dell became a Fortune 500 company. We were doubling revenue. At least doubling every year, it was that kind of thing. There were 700-ish employees when I was hired. The problem is I'm a technical person. I ended up within a year managing more than half of the company. I may be a technical guy but I've been a manager at IBM for 21 years-- 19 years, actually. And IBM trains its managers well. Right? I hate management but <laughs>--

Krewell: You're well-trained.

Henry: -- I knew how to do it. So, in fact, I remember one day, Michael called me over to his house. Remember he only lives a house away.

Krewell: Yeah.

Henry: And said, "I want you to take over manufacturing." I said, "I don't know anything about manufacturing. <laughs> I don't Michael, know anything." And Michael, "You're a manager. Managers manage. You now own manufacturing." So very quickly I ended up not only with the R&D group but manufacturing, I.S., technical support, facilities, basically everything except the sales and marketing groups. And I hated my job. Michael knew I hated it, but so be it. So, Dell grew and grew and grew and after 2 or 3 years, I was able to start getting rid of management jobs, got rid of everything, manufacturing, procurement and R&D, later worked myself down to R&D. And then Michael and I basically, I mean I was really hating my job, Michael and I basically came to an understanding that I would become CTO and stay for a while and help hire my replacement, which I did. But in the fall of 2003, I was on my way out. But I worked for Michael for 5 years and learned an infinite amount and I give him some credit for some of the things we do, and credit for some of the things I refused to do <laughs>, because we're not run like Dell. So, in the meantime, I'm CEO-- I mean I'm--

Krewell: CTO.

Henry: -- CTO of Dell and I have Terry working for me, my old guy from our fellow group, a processor man. And we would sit around saying, "Boy, if we could design a low-cost x86. Look at that thing. Intel is selling something $150, $180, something like that was the lowest priced thing they would sell at Dell." And we'd look at it and say, "You could build that for 10, 15, whatever." So, we wanted to build an X86 but didn't know how. And one day, I don't remember the details, oh yeah, how we got connected. But a person that had worked for me at IBM, in fact, I think I hired him into IBM, Tom Whiteside was the president of MIPS at this time. MIPS had-- was a company that build MIPS computers but they were owned by Silicon Graphics.

Krewell: Right.

Henry: So, I ran into Mr. Whiteside and we got to talking and I gave him my story about, I said, "You've got a RISC-like thing there sort of in MIPS. But I know what Intel is doing and Intel is going to take that ugly architecture and they're going to make it fast <laughs>." And they've got a million people to do and you've got 200. And so you should think about that." Anyway, he hired me as a consultant. So, I left Dell. Had a job and my job is to be a consultant to MIPS to help them develop a strategy for the PC marketplace. At this time, Microsoft actually had a version of Windows NT that ran on to MIPS.

Krewell: Right. And that was pushed by December, was it that kind of MIPS workstation at one time.

Henry: I don't know. Anyway, a week later, Terry quits Dell to-- and I got him a job as a consultant. And a month later, Darius quits Dell and I got him a job as a consultant, too. And there was another person that had worked for me at Dell, Al Sato, and he quit. So there's four of us, we're working as consultants, working out of our homes. I would travel to Silicon Graphics almost every week but the other guys rarely would travel. And we were preparing, we wrote specs, we wrote-- we did design, we were preparing a strategy for Silicon Graphics to compete with the growing PC capability. And what we found is that the Silicon Graphics executives were unbelievably shortsighted, egotistical, whatever. They said, "Ah, a PC, that will never threaten us. Intel, they can never build anything as good as our stuff." Right? This is 5 years before they went out of business <laughs> and 4 years before they converted to Pentiums or something like that. They just wanted no part of this.

Krewell: It was Itanium actually--

Henry: Okay. But anyway, so but they were good about it. They said, "Well, you've got this idea at home." I'll describe the idea in the second. "If you can sell this idea to some of our partners," because MIPS had all these partner people building MIPS under license, "If you can sell the idea to some of our partners that's good so you just keep-- go ahead and do your idea." So, they let us run but it was up to us, but they weren't going to help us. Our idea at that point was surprise its amazing <laughs>, a set of instructions to go into a MIPS that would execute x86 as well as MIPS. Right? And we had a spec and we really

understood x86 by-- at this point. By the way, when I was an IBM fellow, I was the IBM official architecture representative on x86 to Intel. I used to talk to, what are the two guys?

Krewell: Crawford.

Henry: Crawford and whatever name, Gelsinger.

Krewell: Gelsinger.

Henry: Crawford and Gelsinger, so I know both of them. I would call them and say, "You're spec says the instruction does this; we don't think it does that." They'd say, "Oh yeah, doesn't do that." <laughs> That's the way it goes. Anyway, so we kept developing our idea, and I started going around to all of the-- not all. I went to partners of MIPS to sell this idea. And all but one of the partners laughed. I remember being in Japan. I won't name their name, but they're a MIPS <laughs>-- they were doing MIPS processors in a big Japanese company, presenting my ideas and they laughed. And Japanese never laugh at you publicly <laughs>.

So, I was really mad because beyond the technical thing, I now decided that we had to make a company for couple reasons. So, I was giving the argument, "We want to create a subsidiary to do this. I want $15-ish million. I want you to leave me alone, no bureaucracy, no one to quarrel, watch me and trust me. We'll have a processor for you in 2 years." Right? And everyone laughed.

And anyway, I met Len Perham, who was the CEO of IDT and they licensed MIPS. They had a 4600 and a 4700 done by a separate company called QED. But whatever, they were sold under IDT name. And I gave this idea to him and he said, "Well, you know, it's interesting." And he said he would send his technical guy, Al Huggins, down to check it out. So, his technical guy comes, Al Huggins, a really smart guy, turned out to be a really good guy. But he comes and he's one of these guys that comes and says, "This is all shit" <laughs>. "You know, ought to be arrested for proposing. Yeah, what is this thing here?" and, of course, I know what I'm doing. It sounds terrible to say but I do know what I'm doing on this. So, I answer, "Well, this is how it works and we can do this. And you got a concern there?" "Yeah." This went on for months, literally. This is in the fall of 1993. I think I've got my dates screwed up here. I'm sorry, the fall of 1994, yes. I'd quit Dell in the spring- January of '94. Now it's the fall of '94, literally went on for months. He'd raise objections, we would go fix them. So, more objections, he's just kicking the tires and he's good at it. In the meantime, Len takes me in front of his board of directors and one of the board of directors, I can't remember his name, is a famous processor guy who did the Z8 is a-- do you remember--

Krewell: The ex-Intel guy.

Henry: The ex-Intel guy.

Krewell: Yeah.

Henry: Yes. Yes, it was an ex-Intel guy. <laughs> And so we're seeing if we can do this thing, and we will be compatible with Intel. And he didn't believe it because he was an ex-Intel guy and just on and on. But they never threw me out. I went to see the board of directors several times with this idea. Now, secretly-- not secretly but when we were going through this, we would have talks that went like this. Someone would say, "Now, Glenn, what's the value of those MIPS' instructions? Why don't you just replace MIPS ADD with a x86 ADD, right?" I said, "God, that would actually be easier to do but I started doing this to help the MIPS people. And so, you know, that's why they're there, right?" Finally, in late March of 1995, there's the definitive board meeting, supposedly the last of its type. I wasn't invited in. I was there sitting in a room waiting and Len came out and said that they had agreed to do this and fund a subsidiary company, and let's try this thing. Len is definitely a visionary, deserves infinite amount of respect and thanks from all Centaurions for this, but he was a visionary. And Al Huggins helped us a lot to get started, he later faded away. So here we are on our official birthday is April 1st, I think we got our first check 2 days before or something like that. Here we are, 4 of us in our homes, no equipment, no offices, and we get out first check and we got a start and that was the start of Centaur.

Krewell: And the name Centaur means 2-headed- 2- animal with 2 <inaudible>--

Henry: That's it.

Krewell: -- a body and--

Henry: Well, it's a body- body of a man-- I'm sorry. The top is a man, the body of a horse and so-- but some times it's called, you know, a man and a monster, something like that. So, our take was, well, the nice man is clearly a MIPS still <laughs> and that horrible monster <laughs> is x86, right? That was it, it's a 2-headed thing however you--

Krewell: <inaudible> but yeah.

Henry: Have it--

Krewell: But at that point in time you pretty much deprecated the MIPS part of the design.

Henry: Official, the hour we started we were doing-- we were going to add x86 instructions to MIPS. But quite bluntly, what was in our head and what was on our design quadrille tablets was a full x86. It took us a month or 2 months to officially change within IDT but Len had agreed to it, right, we all knew what we wanted to build. So, I was also a vice president of IDT and president of Centaur, so off we started. And, again, they left us alone. They sent money and left us alone which is what I asked for. I mean that's all you can look for in a manager. Right? <laughs>?

Krewell: Yeah.

Henry: That's a perfect thing. And I would go out there every couple of weeks to give status reports, talk about what we're doing. And we even had a marketing group, Centaur Marketing Group, we had a marketing strategy. And the technology was IDT technology. They had a fab there, I think 32 microns, something like that.

Krewell: Sure, <inaudible>.

Henry: Very old <laughs>. Anyway--

Krewell: You had 2 years roughly to get those designs?

Henry: Okay, well, we had 1 year. So, something I haven't mentioned. I didn't mention it to the people right away is the board of directors said you should be able to tape out a chip in one year that can run DOS and some Windows. So, it took me about a month to let people know that though, so we had a year and we did. We taped out a part in 13 months after we got our first check and it ran DOS and some Windows. And the second version we did was the part that shipped. So we did- we sampled a part 2 years from when we started and we shipped 3 or 4 months later. I remember going to the Microprocessor Forum, this is in the fall of '97, where I introduced it. In the movie I gave you there's actually a clip of me introducing at the Fall Microprocessor Forum of 1997, and it was called the IDT WinChip. It had a name. It had branding. People sold it. We had a, you know, a marketing group and we ended up selling several million of those things.

Krewell: And plugged into the 486 socket or the Pentium socket?

Henry: Pentium.

Krewell: Pentium socket, yeah.

Henry: Yes. It was plug compatible. In fact, somewhere in my office I still have a box from a vendor that sold a plug-in, used our part as a plug-in <laughs>, you know, for Intel. Yes, plug compatible, perfectly software compatible. We figured no one was going to do anything software uniquely for us so we had to be nauseously compatible. I think in the history of lifetime, Microsoft has written one line of code or something for us. We are just nauseously compatible with Intel because we have no choice; we have to be able to run code that was done for Intel. AMD more software research, have more market shares so they get more things done for them. By the way, at the time we started, AMD, the only part they had with x86 was exact copy, sliced down copy of the Intel part. And they and Intel were in an argument over a microcode at the time.

Krewell: Right, that's the 486 microprocessor.

Henry: Yes. The- I don't know if it was 486 or Pentium--

Krewell: Was it never built? AMD never built any of the <inaudible> K5.

Henry: Okay, but they were going to-- but they were building the K5 and I met Mike Johnson <laughs>. When we're getting-- when we're in 6-month period we were selling, we met with Mike Johnson and another guy and they were doing the K5. And we sort of told them our idea and they told us their idea. And I remember-- this is easy in retrospective, it's true. I remember asking what their benchmarks were and their benchmarks were bad. And I remember saying, "That's not a good benchmark for x86. <laughs> You guys, there are calls, there are call gates, and all kinds of horrible things and they actually show up in code." So, I think we had-- they were definitely good engineers but they were coming from the RISC, what is it, 2900?

Krewell: 29000.

Henry: 29000, yeah, which was a pure RISC thing. We were coming- we had been PC guys for a long time <laughs> all the way back to IBM, certainly and Dell, and we knew what the software was like. And this Windows software is not efficient, at least in those days it wasn't efficient and it did all these unnatural, weird instructions and things, you just had to support them.

So later, the K5 came out and it had performance problems, the benchmark problems running the benchmarks of PCs.

So, we're up to '97; we have a part we're selling. That first part sold a few million, not many. By Intel standards we're below the rounding error. <laughs>. Okay, but I haven't mentioned one reason why I wanted to start a company - I passed over. Well, there's 3 things about starting Centaur. There's the

business case. We had a business case: at the time I was going around selling it, 50 percent of the PCs in the world were sold in the United States with 5 percent of the world population. So, my story was everyone would want to have a PC, would want to have some form of computing - so human need. And that most of the PCs in the space of time are going to be sold outside the United States, and people don't have a lot of money, right? They certainly don't have the disposable income we have in the United States, where we all buy the best set of tools we can. And so, we had a business story that said we're going to make something that's perfectly compatible, because software is the key, and it'll be lower cost. Later we added lower power, but quite honestly at first power wasn't important then, it was going to be lower cost and, say, therefore it will appeal to the 95 percent <laughs> who aren't buying those Intel things today. So that was our business case but that wasn't why we started the company; that was a justification. We started the company for 2 reasons. One was Terry and I wanted to build an x86. We're processor guys and x86 was built by Intel to be this thing that you had to have a special sauce to build. No one but the finest wizards at Intel could do it. Remember this was a period that AMD had not yet shipped their own x86 design that they finally, you know. So, we wanted to build it because it was a mountain, it was there, it had to be climbed. But the other reason is, I'd been managing there for a long time at this point and I wanted to run a company-- a technical company the way I thought it ought to be run. IBM was a great company when I was there but they did a lot of crazy things, and I'm-- so in my mind there were the good things and the bad things. Dell, for what it was, was a very good company but they did a lot of crazy things, right? So I had the idea of how I wanted to run a technical group. I had- there were bullet points on the chart and things like that. So we wanted to build our x86. I wanted to run a company the way I wanted it to run. And we had a reason to be, a low cost x86, x86 market is only going to grow, low cost. Oh, let me have a drink of water here.

Krewell: Sure.

Henry: So, we build our first part. Then we built a second part. And we were working on our third part in 1999 when Len calls and says, "The board of directors has decided they don't want to be in the PC business." And we weren't doing bad. We weren't doing great. Our parts were selling, you know, but "board of directors doesn't want to be in the PC business. So, we're going to sell Centaur." And he says, "We're going to take so many months to sell it. And then if you can't get sold, well, that's the end of you." So, I actually went and told the people what was happening and gave them the-- those either like it here, don't give up <laughs> I think we got a good story but you ought to know that we're going to get sold. Right?

Krewell: Right.

Henry: There's a risk in that. So, during the summer of '99, several people came to do due diligence. And at this point I just realized I've been rambling for hours and I'm not answering your questions. So--

Krewell: That's okay. There's some times when I'm talking.

Henry: Okay. Should I keep rambling?

Krewell: Yes.

Henry: All right. So, people came to due diligence-- and due diligence, they would bring 3 accountants, 2 lawyers, and maybe someone who knew something technically. Right? And I had a- I still have it. I have a pitch book this thick. You know, all the foils in here that lasted a whole day, the whole story; here's who we are and here's what we can do, here's our patents, here's our talents, here's our tools, here's the number of chairs <laughs> we have, drinks in the freezer, everything a due diligence team would want to know. Right? And the technical guy would talk to me and the lawyers and accountants would be off looking at the books or whatever. And they came and nothing happened. One of them was a Taiwanese company, I can't remember his name. But anyway, one day, Wen-Chi Chen, chairman- the CEO of VIA, a Taiwanese company, calls me on my phone and says, "I hear you're being sold. Would you be interested in being bought by me?" I said, "Sure. Come on by." He says, "I'll be there in 2 days." Oh, okay, so I'm preparing. I got the pitch book expecting a crowd to people. So, Wen-Chi comes with this one more- one other person who said a whole word during the whole thing, I tell you about him in a second. So, there's just two people. They come sit in the conference room. I start my presentation. My presentation is all philosophy, small is beautiful, small team can move quicker, can do more, the productivity is great, no overhead, blah-blah-blah. Right? <laughs> Thirty minutes into it, Wen-Chi says, "This is great. I believe that stuff. I'm with you. I want to buy you. How much should I offer?" <laughs>

Krewell: <laughs>

Henry: Okay. So--

Krewell: Well, and he already just bought Cyrix or is this--

Henry: Well, I-- yes, but I didn't know that at the time. Okay, yes. I didn't know that at the-- he never mentioned that. It turned out that the-- and we talked a little bit more and he went off and he called Len Perham and I believe they consummated the deal in 2 days, it was that straightforward. Or maybe at most a week, it was consummated very quickly. The other guy was the older brother of Cher, Cher Wang who's Wen-Chi's wife and one of the richest women in China, who's from--

Krewell: Taiwan.

Henry: -- Taiwan and China, yes, <inaudible>. Her father was the founder of a conglomerate, Formosa Plastics. He's got hospitals all- you know, lots of things. So, I assumed he was the money man <laughs> but I don't really know, he never said a word. So soon thereafter the deal isn't final yet and Wen-Chi calls

me and says they've bought Cyrix and when I go up to meet he and Cher in Dallas and do a review of Cyrix. This is like September-ish, I guess. So, I go up and, you know, they have a review and people come and talk. And they were about to release a new part and their part had a problem, had a virtual addressing problem with it. And so I asked, so you have this problem, yeah? So, I ask to tell me about what they knew about the problem, what they were going to do about it, et cetera. I didn't like their answer. So anyway, at the end of the day, Cher and Wen-Chi are in the office and they ask me what I think. And I say, "That part will never ship. You just do what you will with that. That part is not going to ship for these reasons. In fact, you know, young cocky guy then, I said, "I'll bet you" <laughs> "on a steak dinner or something, that part will never ship." And he just didn't make any response. There were also lots of personnel problems. The previous owner had done crazy things when they were going to sell them, we won't go through all that. The upshot is that Cyrix within 2 months swindled to like 10 people or something. Basically they, the whole company disappeared. How much of that was them dwindling on their own, or how much of that was Wen-Chi, I never knew. But I do know- I mean I had to be parochial, that's my job here. I do know that I remember presenting Wen-Chi a chart after a review that said here's how many people they had. They had 450. They had 15 people with titles of directors or above, right? And at the time, I had 60. Centaur was 60 people. I said, "Here's the story, 60 people, 1 manager, everyone else is a worker where it's just 400 and something." So, I made that point strongly because he asked me if I wanted to hire those people. I said no.

Krewell: Well, now, Cyrix was owned by National, I think at the time, and they were being sold and National is a subsidiary.

Henry: Yeah. They had bought it and, you know. They’re a smokestack -- I shouldn't say. They're not really a smokestack <laughs>. But they had fattened it. Cyrix was lean and they put in vice presidents and directors, a typical thing a big company does. And remember, Wen-Chi inherently likes the small, lean thing. And so, we were small, lean at 60 people. They were 400 people, the handwriting was on the wall but I don't know the details, but Cyrix sort of evaporated. And that part never shipped. What they did was take the part that we had been developing, because we were still developing. I kept this team alive. One person-- one thing I'm very proud of doing in that period, 3 month period, maybe 4-month that we had announced we were going to be sold, one person left. We just kept working on the part. So, when we were bought we had our part just ready to tape out. <laughs> VIA sold that as a Cyrix, put the Cyrix name on it because they thought they had got a brand in Cyrix but it was our part.

Krewell: Yeah.

Henry: I forget--

Krewell: There were very limited number x86 competitors at that time--

Henry: Mm-hm.

Krewell: -- and competing with Intel aimed to use large <inaudible>. Cyrix was relatively well-known and find <inaudible>.

Henry: Yeah.

Krewell: And then there's a few other older clones like the IBM Blue Lightning-- <overlapping>

Henry: Right. And there were all kinds of people trying to do it. Chips and Technologies, an old company tried to do it. Lots of companies and startups were on those days were trying to do it. Now I think people <laughs> be crazy to try to do it. But in those days-- because we met with people at conferences, the Microprocessor Report too, all kinds of people that would come up to me and talk about doing it. And they all had great stories; how fast it was and, of course, most of the time I did. But every once in a while I'd say, "So how are you going to handle the bounds instruction?" "Huh?" <laughs> You know.

Krewell: And there's a bunch of oddball instructions and that <inaudible> you have to take to account in the next <inaudible>.

Henry: Yeah. Anyway, the-- we had a part taped out that fall and shipped the following spring. Things worked fast those day. It was labeled to Cyrix but after that they didn't use the Cyrix name anymore. All the parts that VIA shipped were ours. On my keychain here, except I didn't bring my keychain, we have a little thing when we sold 10 million parts. It was soon after that. It was with VIA but it wasn't that much longer, we had a little keychain made with a chip in it. So, we were selling a few million a year during that timeframe.

Krewell: Is part of the reason you could build a chip so great was because you went with a very simplified architecture? You kept it as minimal as possible, kept the cost down? It also made the design easier to produce?

Henry: Yes. Small, all good things come from small low cost, <laughter> quick turnarounds, small teams. Small team further begets quick turnarounds, right, and low cost, et cetera, et cetera. Going back to the business plan, I'd seen first-hand, and participated at Dell, in the problems of growth, you know. And so, I said, "Rightfully or wrongfully we are just going to be a design company. If our part will sell billions, that’s fine. The sales and marketing, the things will be somewhere else, right? We are designers here and that we're going to be so inexpensive doing our design that the chip just has to sell Intel rounding numbers to be successful." Right? So, we live a good life here, <laughs> and I have good equipment but we're small. Years and years later, now we're 103 people today, right, and we've been roughly at 100 for the last 10

years. My budget hasn't changed more than 20% what I spend in the last 10 years. So, we don't-- compared to the tech industry which you're very familiar with, we spend nothing <laughs> and that allows us to stay in business. Now, we're not going to take over the world but my theory of this business, if I may deviate, is threefold.

We want to do good technical work, technical work that's hard and we want to do it well with quality, and X86 is a very good test of that. Secondly, we want to have a good life. People here like-- say that they do. They like working since our turnover's negligible. Right? You can find that out lots of different ways. Our turnover is very low. People come here and they stay here. Of the original 25 people we hired at the end of the first year, 18 are still here. So, we've provided 22 years now a good life for the people that we hire, and third of all, we want to try to make a difference. Our difference hasn't been nearly as high as I'd hoped, but we have shipped tens of millions of parts. We've stayed in business for 22 years and the parts were truly lower cost than Intel.

The tradeoff was twofold. One is the parts weren't as fast as Intel but they weren't as costly, partly because they weren't as fast, partly because we did it with 26 people the first year, then 50, then 60, and then 100 at peak time, and that money has to be accounted for somewhere. The speed thing has always been a problem for us. You've heard me, I think. For seven or eight years I would go to all the conferences. I think I talked at Microprocessor Forum six or seven years and the story I was always telling was you don't need to get a faster PC every two years. You don't need a new chip. Well maybe you do, you audience people, but ordinary people don't. What ordinary people do, they browse the internet, they do email, they do spreadsheets, what have you. So, I used to give a story in my talk, if I can ramble. I used to say I know in my life there are four grades of PCs. At Centaur we buy the fastest, biggest thing Intel can make. We have a building with thousands of them. Okay. Then I said, "my son buys the fastest thing because my son was making a movie at the time, right?" But not quite as fast as our servers. He runs a fast desktop. I said I have one level down because we don't do development on my machine. My machine's a management machine. Alright? My wife's using a Centaur chip, alright, because what does she need better than that, because here's what she does, right? I said so there's that spectrum in the world, but down here where my wife is, that's 95% of the world, right, and that's what-- well maybe it's 90%. Who cares? That's what we're targeting.

Krewell: Right.

Henry: Ordinary people doing ordinary things, and that's why the PC is so valuable. That's where all the volume comes from. And so, we built a part for those people and yes, we're not as fast as the latest Intel chip. Wait two years and our today part will be as fast as <laughs> today's Intel's chip. You know, we're two years behind them, right, and it's fast enough to do the job, but that never really worked, Kevin. I used to-- first of all no one likes to hear me say we're fast enough. That's not good marketing. <laughs>

Krewell: Well <inaudible> talked about good enough PCs for a long time.

Henry: Yes, but that is our story and today I think it's even more clear, especially with the price of some of the-- you can spend a lot of money on a processor these days and they're very, very powerful but not everyone needs one.

Krewell: Well Intel even responded with their Atom processor which is also a stripped-down core, design core, more entry level and, you know, fundamentally.

Henry: Right. Right. It doesn't fit. I think the Atom was a good part. They've sort of giving up on it though because well the problem is, of course we all know this. Intel's got a 50% gross margin and a massive R&D expense, right?

Krewell: Yeah.

Henry: You can't afford to sell those kind of parts. Our strategy was if you're small you can build a part that's low cost and do it quickly, and because it's low cost <laughs> you can afford to do it small, and Intel's the opposite. They've got to make a lot of money on every part they sell.

Krewell: Right. So where do you sell most of your parts today? Is it mostly the rest of the world?

Henry: Yes. It's all-- well, it's not-- over time HP's been a good customer of ours. Primarily they did notebook but primarily terminals they've had. No, what's the name?

Krewell: Terminals? No?

Henry: No. They're terminal-- they're like PCs that have sort of no guts, thin client. I'm sorry. Basically, we're going back to the IBM days <laughs> with terminals to mainframes except mainframes we call them servers now. Anyway, but now all of our stuff is primarily overseas.

Krewell: Mm-hm.

Henry: Lenovo actually, for example, sells our part today, but they sell it only in China on a Chinese SKU. That was always sort of my strategy. Anyway, yes, that's where most of our sales are.

Krewell: Okay. So, you've been competing with Intel or in a market with Intel for a long time and the only other vendor that's been competing with them longer has been AMD. So, do you have any interesting stories about even Intel at that time?

Henry: Well, I can't tell you the marketing and sales stories. I can tell you apocrypha from the marketing and sales guys <laughs> but I can tell you the story I was involved in, which is the patent feud of course.

Krewell: Yeah.

Henry: Should I tell you about patents?

Krewell: Patents are interesting.

Henry: <laughs> Maybe. I hate them. So, when we started, we started doing patents. I mean Terry and I both have a good background in patents.

Krewell: Oh, IBM's..

Henry: IBM and Dell.

Krewell: Yeah.

Henry: Terry was the force that got Dell out from under TI patents. He spent over a year doing that at Dell, but anyway because in this business and tech business in general you've got to have patents. In particular, Intel has sued everyone who ever did-- basically we talked about Cyrix. They sued Cyrix before Cyrix even shipped one of their chips..

Krewell: Yeah.

Henry: ...for patent infringement. For various reasons, and we have our theories but we don't know, Intel did not sue us for several years. So, we just kept selling parts, and what Intel did in 2001 was to sue VIA for VIA chipsets. At that time VIA had a great market in Intel compatible chipsets because Intel had gone to another kind of DRAM or I forgot.

Krewell: It was RDRAM.

Henry: Yes.

Krewell: It was RAMBUS.

Henry: Yes, RAMBUS, exactly where VIA was making chips with good old DRAM, right, and RAMBUS had problems. So, Intel was doing very well. So..

Krewell: VIA was doing very well.

Henry: VIA was doing very well. Sorry.

Krewell: So in terms of just <inaudible> Intel processing.

Henry: Yeah. So, Intel sued VIA over the bus patents and VIA has effectively no patents, no way to answer that. So, I remember getting a call <laughs> saying "We've got to look at your patents and someone's coming from Wilson Sonsini to talk to you." So pretty soon there's a group of lawyers in our conference room and we're going through our patents. We'd only been in business for four years now. And so, we found two or three processor patents that we thought the Pentium 4 infringed. So we sued, we, Centaur and VIA sued Intel for..

Krewell: <laughs> <inaudible>.

Henry: The day later they countersued. So clearly they were ready. Right? It's interesting if you go and look at the material, there's a picture of that in this movie I gave you, the suit was filed on September 10, 2001.

Krewell: Oh, God.

Henry: Yes, the day before <laughs> 9/11. Yes, and our lawyers were supposed to come down from New York the next day. It took them a week to..

Krewell: Yeah.

Henry: ...make it here. So, a memorable occasion in two different ways. So, for a year, almost two years I spent at least half my time on that suit. First of all, I was president of the company. Second of all, all three patents we used I was an inventor on. Third of all, our lawyers were able to get me certified as an expert witness even though I was a party to the trial.

Krewell: Yeah.

Henry: Intel protested that I was an expert witness, and patent suits have many steps before the judge, before you get to a trial with a trial, but because we had sued-- so I had many days of depositions and I was testifying in front of the judge two or three times before we even got to the dang trial. Since we had sued them first we asked for a jury trial here in Austin. So we had venue..

Krewell: Okay.

Henry: ...jury trial, and for the week before the trial, the trial was supposed to start on my son's birthday, April 7, 2003, I had spent the week locked in a room down at the Four Seasons with the lawyers because the entire first day of testimony was me because I was presenting the story of Centaur. We are just a bunch of good old boys trying to lower the cost of computing for you farmers and the jury <laughs> and here's all of the wonderful stuff we've done. Here's my boy, Terry, down there. Terry's wearing an ill-- he would come to all of the trial things, too. So, he wears an ill-fitting suit and he looks like a typical engineer. There's that California corporation and there's their New York lawyers.

Krewell: Yeah.

Henry: We had New York lawyers too but we had a front man who's the local <laughs> friend of the judge..

Krewell: Oh wow.

Henry: ...who would get up and say "Yeah <laughs> I'd like to explain these patents but I want to let these guys do it," you know, and they're coming first. Anyway, I spent a lot of time on that trial and we were ready to win. I was putting on my tie which I hadn't worn forever. That morning I get a call from Wen-Chi saying that they had worked all weekend and they had settled all suits with Intel and the suit was a cross-license agreement with all Intel patents. So to me that is a huge win. You couldn't buy an Intel cross

Krewell: Right.

Henry: ...license for any amount of money and we got one and it was worldwide, too. So, Intel, of course they had done the right thing with the counter suit as they counter sued us in 18 different countries, right? In fact, I had to send Terry to England for two weeks to be at the English preliminary trial, right, so we had someone that knew something there. I think the way it was going to come down, you know, you get a pretty good feel from the lawyers of the strength of the case, but they were going to win on one of their patents, on MMX. They have an MMX patent. <laughs> You know, if you do MMX it looks pretty clear.

Krewell: Yes.

Henry: And we were going to win on one of our patents which was a new instruction that they'd put into the Pentium 4..

Krewell: Mm-hm.

Henry: ...that we had invented because they weren't doing MMX. In fact, it goes back to Microprocessor Forum.

Krewell: I was reading that.

Henry: The-- yeah. Well my version is through them. So, in 1997 at the Processor Forum, or was it '98? I can't remember. AMD stands up to talk about their 3D Now and I stand up and present our version. Alright? And Intel's dead silent, right? Intel's taking a lot of heat because they don't have..

Krewell: Right.

Henry: ...a 3D Now in..

Krewell: <inaudible> said it was beyond just <inaudible> to California.

Henry: Yes. <inaudible> So anyway, and somewhere at that meeting, I don't know if it was public or it was in a private conversation, Microsoft representative <inaudible> says "Intel's going to do something, and we're only going to do one other-- support one other thing," and the Centaur <laughs> one other thing, you get it?

Krewell: Yeah.

Henry: So I and AMD reached an agreement very quickly that we would license their 3D Now and solute it,<inaudible> right?

Krewell: Yeah.

Henry: Put their trademark on it, what have you, and they would get a license to our patents related to our technology. Anyway, one of the patents that was involved in that, one of our patents is this instruction that Intel probably infringed, but who knows? The suit and many things going into suit. Later there was one other suit if I haven't mentioned which I really hated was Apple sued HTC for infringement on their cellphones. You know, HTC cellphones?

Krewell: Okay.

Henry: Okay. Well HTC is related to VIA. Cher, our Chairman of the Board, the CEO of HTC when she's on the board of HTC, alright?

Krewell: So then Intel <inaudible>?

Henry: Yes. I don't know what the Japanese word is keiretsu. I don't know what the Taiwanese word is but they're interlocking companies, right? And so, once again they come to us for patents to attack Apple. So, we sued Apple back and the ITC or International Trade Commission, and I went up there and spent two weeks in New York with Terry again, and I testified for almost a full day because the number one patent we had to assert was the same patent we had in the Intel case.

Krewell: Oh, wow.

Henry: And it was very painful. I mean their lawyers were very good. They'll make you think you're an idiot in 30 seconds, but anyway that case never went anywhere, because it was settled very soon after that. So, we've had-- our patents have been used in two cases that resulted in settlements with Intel and Apple, or two of the most powerful entities in the world. That is a segue of I hate patents. I have over 250 U.S. patents. I still hate them, but it's a necessary life.

Krewell: Yes.

Henry: And we're pretty good at it here.

Krewell: Okay.

Henry: Ask a question, Kevin.

Krewell: Alright. We've covered a lot of what's happened over the years. Your career from IBM produced Dell and side projects and Centaur. During all that time what do you think were the high points of your career?

Henry: Well, I think it's coming up with the original technical ideas and causing the System 38 to be developed. It was very, very hard. It was hard to do just technically. With the technology in that day, it was very hard within IBM. I mean I fought and pushed and lied and <laughs> whatever you do to get a project done, but technically it was difficult. We were doing 48-bit addressing, single global stores in the mid-1970s, right?

Krewell: Yeah.

Henry: Two size page tables, two size of pages of stuff like that. That I think is the technical high point. Certainly from a managerial viewpoint, creating and managing and keeping Centaur alive for 22 years is the high point.

Krewell: Yeah.

Henry: Managerial way.

Krewell: Would you say that's been the most fun and rewarding has been the Centaur or..

Henry: Yeah.

Krewell: ...or System 38?

Henry: Well it depends. System 38's-- a lot-- probably no one knows what it is. I know what it is and I know what became of it. So that's something. It's just a personal thing. It was hard to do. We did it. Centaur's different. It's hard to do in a way. It's more the feeling, the positive feeling that come from having a company that's survived for 22 years, employees never leave. They like it. Over that time, in fact I don't know why Obama or Trump never give us an award for this because we're the perfect-- what we have done is bring foreign money into-- all of our money comes from Taiwan, really China these days effectively.

Krewell: Yeah.

Henry: We brought all the money in to provide good jobs here. In our lifetime here, we have spent over $500 million here in Austin. We buy from Dell, building, people's salaries. So, we're bringing in money to pay for Americans here. But, forget that. Providing employment for 100 people who like it and it's good employment. It's a good job. We take care of people. That's certainly a great feeling.

Krewell: Okay. Well..

Henry: So in one sense I guess that is. I have to rate that as my best accomplishment. The other is more personal <laughs> versus sort of..

Krewell: Well it points to the-- in one sense the importance the intellectual property. You mentioned how you're not happy about the patents but patents are extremely important as part of the intellectual property that you've generated.

Henry: Right. We do it and we do it well. One of the things-- we learned a lot. I'll mention it a little bit. In the Intel trial, Intel had thousands of patents. They could find five to assert against us, right, and two of them were thrown out by the judge before we got to trial. We asserted three. One was thrown out. One of the things we learned is that I don't want to say it's easy to get patent but it's totally different than getting a patent that survives litigation. It's relatively easy to get a patent granted. It's very hard to get to win in a patent litigation. And so, what we have done is we focus on good patents. By good patents, we don't do patents for their sake. I have, as I say, over 250. I could have a lot more if I just wanted to do patents. We do patents that we think will have value and the value has to come from the potential of litigation. Right? Patents that can survive..

Krewell: Right.

Henry: ...in a court case. So yes, I shouldn't say I hate patents. I just don't like going to patent trials. <laughs> I think patents are useful. They're important and you have to do them.

Krewell: Alright. So, what are you working on right now?

Henry: Well, I'll give you two answers, what the company's working on and what I'm working on because when we got to 20 years two things happened. We had a great party. <laughs> We have film of it, fireworks, entertainment, you name it. Twenty years for a crazy idea. I decided it was time for me to stop doing everything I was doing. I'm the only real manager here, and I was driving technical things. I was writing-- I've written microcode on all of our projects. I did hardware on one but mainly microcode because..

Krewell: Mm-hm.

Henry: So I have released, I'll call it what you think of it, the project management type stuff I was doing. I still manage people and money but I started-- I appointed myself the research division of Centaur. I was <inaudible>..

Krewell: Appointing yourself as your own CTO, right?

Henry: <laughs> Yes, my own CTO of my own research division. I started a project. It's something I'm really interested in, but it's actually coming out. It's going to have merit because along the way I got two graduate students from UT and two other guys contracted from another company. One guy's a world-class person doing the software and I'm building the hardware for something that I think is going to be quite useful, but it certainly is a lot of fun and it's a co-processor to our X86 processor that has nothing to do with classical computing. That's what's the modern trend now and some degree is domain specific architectures. It's the domain specific architecture that I'm very interested in, and I was all the way back in IBM. I always used to believe that-- I mean if you want to ramble about processors now that the rate of progress is very narrow. If you think about the past, first there were just the basic architecture things and pipeline and then super pipeline and then there's out of order execution, right? And then all along the way Moore's Law's working fine, and then there's multicore, alright, and now the main thing that our part that our team is doing is an eight-core part. We're in the low end of the marketplace and we got an 8-core part, 16 meg L3, you know. I mean I think what's happened now is Moore's Law's no longer working as well and the end of the tunnel is in sight. We could make a chip with 100 cores on it and no one would know what to do with it basically. Right? So, I've become a believer that the curve is like this, the processors. Yet if you look at their specific computational intents of things in the world where a specialized architecture to get a 10X increase over our processors where if we put all of our-- the difference between processors now is in tens of percent. Yeah. Yeah. Intel's going to be 20% faster or something like that, you know, the new eight core processor. Megahertz really isn't growing. The power has become limiting. The new technologies are just terrible in terms of wire delays and I'm telling you things that everybody listening to this knows. So, I've become interested in building specific architectures for specific problems. In my case I'm building something. I'm doing all of the hardware. It's a one-person hardware project and yet I'm going to provide a performance benefit for this particular thing of 10X over our eight core X86 processor, and I think that is an interesting trend.

Krewell: So in other words an energy efficient computer model going forward where you have different processors doing..

Henry: Yes.

Krewell: ...different things.

Henry: The thing I'm building is a processor, has an instruction set, very weird but it goes really fast. Something that's very important, a very important thing done in the world today.

Krewell: Alright. So, it sounds like you still have something in the bag that you haven't put out yet.

Henry: Yes. So, I don't know the timing of this, but I shouldn't say more about it, but one could look around and think what things in the world need massive performance and everyone's talking about it and everyone's doing it. There's a hundred startups. You could find a hundred startups <laughs>. Do you know Chris Rowen?

Krewell: Yeah, sure.

Henry: Sure. Well Chris, you know, has got..

Krewell: Former founder of Tensilica.

Henry: Yeah, founder of Tensilica. Now he's going in sort of the venture capital with his own money but he's specializing. He's got a char on his website of hundreds of people working in..

Krewell: Yeah.

Henry: ...this area.

Krewell: Yeah.

Henry: So but so that's what I'm doing. I'm..

Krewell: No. It sounds like a very interesting project and still relevant to what's going on today. So you're still..

Henry: Right. Yeah. There's a lot of people doing what I'm doing today. They're all doing it differently, of course different, but most-- a lot of them, in terms of what people have done so far, we're the first person to marry the specialized computing power with the power of the eight core X86. This is a co-processor. We can mix and match. If you've got some-- most of the calculations of this domain are straightforward and we do them really fast, mass parallelism kind of thing, but every once in a while

there's something that's really weird to do. Instead of doing hardware, it doesn't happen very often but I don't have to do it because I've got high speed DMA into the processor. They can do it.

Krewell: Yeah.

Henry: So it's the fusion of two computing worlds, where each does what it is really good at doing and yet can offload work to the other one.

Krewell: I think it's one of the ones that <inaudible> GPUs and CPUs. It sounds like you've got a lower cost, lower power solution.

Henry: Well, I'd say certainly lower cost. I think everything's lower cost than it what <inaudible> 500 square millimeter chips would have you. That's not our style.

Krewell: Well, (NVIDIA) Volta is an 810 millimeters squared.

Henry: Yeah. I know. I couldn't believe they could make a chip that big.

Krewell: Yeah. Most of them totally do. Yeah.

Henry: So it's been a good career. I've done lots of different things.

Krewell: How would you summarize it?

Henry: Well, it's a great career. I'm very pleased with it. I would point out that I've been incredibly lucky. I was-- not in detail, I was a right time, right place, right person in all these things and I'm very grateful to that opportunity. I mean I won't go through all the details but I will mention some. The opportunity I got at IBM, four promotions in the first four years, very rare. There were other people that are smart there, too. I was in a group. I was allowed to do things especially and I had an understanding of managers. Michael Dell moves one door away from me. Alright? I ran into Len Perham]. I ran into WenChi Chen These are unique people, alright? That's what it takes, but I feel that I've done the best. I feel I've taken advantage of those opportunities and I've started by being excited with the IBM 650. Remember a million years ago..

Krewell: Yeah.

Henry: ...when I was in the 11th grade and seen it grow to what we have today, and that's a good feeling and I still love the technology and there's still more to do.

Krewell: And you still go to work every day, too.

Henry: Yeah.

Krewell: So that's great. I think we've covered the career. Is there any last points? I mean is there any-- do you think if you ever come up to Mountain View, I mean obviously we'd love to have you see the Computer History Museum. I'm not sure if you've been to it yet.

Henry: No, I haven't.

Krewell: You have your own computer history museum.

Henry: Yes, which I want to take you to see. Okay?

Krewell: Yeah, and then if you ever get up to Mountain View, we'd love to give you a tour and show you. You might find some System 38s or other standard..

Henry: Yeah. I'm sure you have things that we worked on, but I haven't really researched it. I don't know. I think the tough-- I'll say one more thing. Today there's a massive amount of computer science and theory and really good work on everything, the languages and processor architecture, silicon technology. I still remember-- this is what you get from old people but I still remember the good old days after we'd walk to work uphill through 10 feet of snow..

<laughter>

Henry: ...there was no-- there was no technology base. We hadn't invented the things we did. I remember going back to System 38. Yeah, virtual memory had been around for like four years at most at IBM. We didn't know how it worked. We were inventing it, right?

Krewell: Yeah.

Henry: Making models. The engineers are dealing with crazy things that you don't have to deal with today, power supplies, water cooling and on and on. To me it's very interesting just seeing it go from a very ad hoc industry to something that's quite scientific now, quite polished.

Krewell: Yeah. Oh, it is. It's a maturing industry but at the same time there's still a lot of room for innovation like this little side project.

Henry: Of course. I mean you can name a whole bunch of things that are important to people, that could be important to people today that need computers 10 times faster. Or more importantly, we probably need data transmission 10 times faster, different bandwidths. So yeah, and the other thing that really-- I don't know what the word is but about my environment. When I started no one had a computer. Well, that's not true. IBM had 10,000 computers a year or something like that, right, and now they are ubiquitous. Half the world has a computer and the other half wants one.

Krewell: And they're popular.

Henry: Yes. So that's a very impressive change to me to see something that I love and something I think contributed to and help work and grow from in the specialized thing to something that ordinary people-- it helps ordinary people's lives. It gives this tremendous value to civilization.

END OF THE INTERVIEW

