BL: We're lookng for in terms of what you remember in terms of significant stories, events, things that happened that give other people a glimpse into, - say, oh, by gosh, that is DEC.

AK: Years ago, 1963 or '4 maybe, I was working at Raytheon, the government had bought as government furnished equipment, ... shipped to Wright Patterson Airforce Base a PDP1 and, part of this thing was, that we needed a lot of analogs to Digital. There was all this stuff in the module handbook in those days but we couldn't quite figure out how to do some of it. PDP1 was almost neutral-like anyway. So we made contact with a woman who may well have been Digital's first woman engineer, probably was, by the name of Barbara Stevenson.

BL: Yes, I'm hoping to visit her in Albuquerque.

AK: She left DEC years ago but she was very helpful. She told us what ADDT converter you needed to buy and showed us how to take DEC modules and build DDA converters out of them. So we did and it worked. Right after we had sent the PDP1 out to Wrighter Patterson, we got a UNIVAC 1218 and by comparison it was a hunk of junk and was probably pretty good for those days but it was sufficiently awful that we took the ideas between ADDT and some of the other stuff on the PDP1 and rewrote them

for a 1218. We were sick and tired in dealing with a 1218's utilities.

BL: DDT was a debugger.

AK: Yeah, DDT was a debugger which was - I don't know if you've run into the story of DDT - I had nothing particularly to do with it. Ann Murphy originally wrote DDT at MIT.

BL: I remember it from another book.

AK: DDT's ancestor was called, - FLIP and ran on the TX0. The PDP1 I worked at Raytheon is an interesting case because it had to go in the triangular corner of a circular room, as it was. The room was just big enough for a human centrifuge and all that was left was the square corners where the centrifuge didn't swing. So this PDP1, we had to bust up into two pieces, and one went on the first floor and one went on the second floor with cables between them and that all worked too.

BL: What was it being used for?

AK: Well, we didn't have the guts, in fact nobody had the guts in those days, to use it to actually control the centifuge. So the centifuge control was all hard wired

and ____ simple besides. Let me be careful. The major access control of the centifuge was hard wired and analoged and the PDP1 was able to read it but it was actually three axes of rotation and three more axes of translatin of the seat where somebody was supposed to sit and three axes of tilt and the small axis of movement, both up-down, crossways of the seat and the tilt axis of the seat were actually controlled by PDP1. Since it was hydraulic, it acutally had limited movement, so it couldn't go terribly far.

BL: Never got out of control, did it?

AK: As far as I know, nobody ever actually rode in the thing, which was not PDP1's fault or Raytheon's fault but rather the _____. It was a frustrating project. I kept track of it for several years after I came to DEC and still hadn't _____. The whole thing ran, in every respect, except the main axes bearings oscillated and you'd start the thing running, just sort of slowly and it started oscillating and you'd stop it from turning and it would continue to oscillate and turn off the hydraulic pumps and turn them back on and it would be stable. The prime contractor, _____ _____ several years after I went to the project. As far as I know they never did get it to run.

BL: [TAPE CUT]. .... originaly designed to be a flight simulator, sort of in the same realm.

AK: What this project was supposed to be was a, called a dynamic escape simulator and was supposed to simulate the forces on a person when they bailed out at Cape Canaveral just after the rocket lifted off. Fortunately, that's never happened in reality. It's supposed to be able to happen. I don't know if it's true of the Apollo. In the days of the Apollo there was this rocket charger that was supposed to blow the entire crew out of the thing and make it parachute down with the forces [INAUDIBLE].

BL: What you mentioned about calling PDP problems, in a sense typical in those days of Digital engineers who were probably development engineers.

AK: Oh yeah, she was very definitely a development engineer. As far as I understand it, she was responsible for all the ADD and ____ products in those days. She was a very bright person.

BL: But that indicates how close the development engineers worked to customers.

AK: Oh yeah, in fact another interesting story which isn't really about engineering, in the same machine, the

same operation Raytheon we had some minor breakdown which isn't surprising. So we called their field service. And, guess who came out to service it? Jack Shields. I mean he might have been head of field service, even in those days, I'm not sure, but he was also one of the two or three field service engineers around here.

BL: Did he get it going?

AK: Oh, yes he got it going, no problem. In that days that was, four or five hundred people worldwide. This must have been '62, or '3, or '4.

BL: What was the outside perception of Digital's engineering, whiz kids, crazy people, tops, mediocre?

AK: I think from the outside it didn't seem much different. In those days DEC's principal competitor, at least in New England, was an operation known as Computer Control Corporation, which has since been sucked into Honeywell. They both seemed to be similar, they built things and they by an large worked, sometimes they didn't and that was the way things were. People would say, well, the NTPF must have been a real _____. PDP1 would typically fail every two or three months probably, which, isn't that bad at all. There's a lot less logic in one than there is in a chip today but still.

BL: How did you then filter into how involved in the mill and what did you find when you got there?

AK: Well, it's complicated. I went to school with _____ _____, and he used to hack the radio stations with me at MIT and he also hacked the model railroad with Kotok. We were pretty good friends, in the summer of 196 he said, - well, why don't you come over and help us work on this new processor and things? And, I interviewed with Bob Savelle, Century Tire, he said, - some day we're going to start up this neat high end special systems group, in customer projects but in the meantime, we got this CPU to get out. Would you like to come help us? I said, - that sound pretty interesting. So I came over in September of 1966 and started doing circuit design for the KA10(?) which, at the time was known as the 166. Kotok and Clemmens were doing the logic design. And we shipped the first couple with basically _____, literally PDP6 peripherals. There was this thing that built, little teeny boards, the late B series modules and the ____ modules, that had the old system modules peripherals in racks(?) next to it. We shipped two or three that way. Then, Clemmons and I and Dave Rose, designed some peripherals for the thing including the marvelous card punch. If you look at the literature when the KA10 first came out, we said, - well

we have _____ and card punch as well as disks and tapes and all the things you got to have. We really did have a wine(?) printer and _____, not a very good one, a tape drive and _____ stuff. We didn't really have a card punch but you needed to have a card punch in order to go into the pseudo EDP shops that we were pushing in towards. The sales people always managed to talk them out of the card punch before anybody actually signed the order. Then, one day, a couple of years later, we managed to get away with this for at least a couple of years, some salesman didn't succeed in talking the customer out of ordering a card punch so we actually had to go out and find the card punch and build an interface for it, ship it which we did. It must have been two or three years, we were selling this mythical card punch. It wasn't as bad as it sounds because we had a card punch picked out. We said, - well, if we ever have to sell one, we'll sell this one. So, it wasn't really blank but it was a little slippery.

BL: What was the feeling about engineering? Could something work? It's not like a too well thought out.

AK: Oh, no, no. I took a lot of pride in designing the stuff and everything we did was done the best _____. Take the DC10 communications controller which is something that Bob Clemmons started and I did most of the

work on and carried through. I think somebody else is now listed as the responsible engineer. That thing was built out of PDP8 single board communication modules that were then mobiplexed(?) eight ways and the only trouble we ever had with that was that the PDP8 design was designed to be bussed together but when you bussed it together it turned out to be difficult to _____. So there were certain _____ difficult to deal with. But we were still shipping that thing into _____, not at very high volume [INAUDIBLE]. The same thing is true of the TD10, I think the TD10 DECTATE controller, was the only DECTATE controller that was ever designed for the tab(?). That thing was designed for the KA tab, was actually sold on the last PDP6. The DECTATE controller, you could buy, if you wanted one at all, ___ __ _____. It's a different environment because you say, - well, we sold this whole thing. Well, I think we sold something around 240 or 450 K810 total, lifetime. And we thought we were doing real good in those days. For those days we were doing good but compared to _____. You talk about now and that's special systems [INAUDIBLE]. We tried real hard to build things as well as we could possibly do. We took pride in doing careful analysis and everything else, even the tools we had. In fact, until the KL10 designed by the PDP10 group was the leader in designing CAD tools for the company. All the tools in the K810, KI10 days were put together by the engineers out of desperation because

nobody else had tools for them. So, I didn't believe in coding the tools but Kotok, Rose, and Clemmons and Edgars, coded almost all the tools and Dick Hallowell in the KI10 days. So all those tools were hand wrought. The only tool that wasn't hand wrought was the original wire wrapped system was a hard system that ____ had put together, I guess, it the PDP8 and we had to expand that considerably but that was the original. The wire wrapped controller software was written _____. [INAUDIBLE] ... six ways to Sunday and at the end of the KI10 stuff, I documented all the software we had to make the KI10 go together. We started calling it kluge forest, because there was forest of kluge towers. Because that wasn't what we were trying to do, we were trying to get something out, it wasn't very well documented and it was perhaps a little fragile in the sense that we tried something we hadn't tried before. That's what we did. That's what we had to do, there wasn't any other choice.

BL: Circuit things were a lot different, [INAUDIBLE] ... tested out.

AK: Oh, it's a lot different. Certainly we bread boarded stuff as we do now. People keep talking about bread boarding and the reality was we essentially never did any bread boarding in the sense of taking a piece of perforated board and soldering transistors together and

seeing how it worked. After we built the standard gate for the KD10 for instance, laid out and built ten or fifteen versions of that gate with different numbers and inputs, directly on to boards that were indistinguishable from production boards. We did build engineering _____ boards they usually were and we had to make minor changes to them but prototyping in the sense of tangles and wires going on, we did very well with that. And, if you think about it, you really can't do that. In those days we were building stuff that was as fast as anybody knew how to go, which by today's standards weren't all that fast. But the KA10 modules had delays that were spec max, 10 nanoseconds, 12 nanoseconds, 15 nanoseconds. And by today's standards, that sound slow but that was faster than the stuff you get in ICs that were available in the 7400 days. They were moving right along so you couldn't just build this kluge of wires and expect it to work like the final board because the final board was laid out, assuming that things weren't just capacitors and resistors, they were really conductants and waves going down the board. [INAUDIBLE]. There were two, Bill Walt(?) was the last of them. He started out as a technician in those days. He's still with the company. [INAUDIBLE]. It wasn't cut and dry. It was, let's understand what the hell is going on and make it work, make it work right. That showed up because those machines didn't have any braking(?) operation in them.

They plugged them in, they worked, they worked _____. If you were clever enough to test the board before you plugged them in, which we were, those things went together like a charm. The first prototype KA10, when we plugged all the boards in, turned the power on, we had a running operating system the next day. I mean, he didn't believe it either. Not only was it a _____ designer, you're damn lucky but it really did. And it wasn't that we didn't have to make a few changes. By the time we ran the operating system, there were 20 or 30 degree(?) wires in the back plane, that was all. We were as amazed anybody was. It wasn't planned that way, that's for damn sure. We couldn't have done it on a [INAUDIBLE].

BL: What were the tools and what was the role of the individual and the role of the team?

AK: I'm trying to remember. As I say the KA10 team was basically Alan Kotok and Bob Clemmons doing logic design. There was me and Bill Walton and Pat Sullivan, that's his name, doing circuit design with a couple of technicians. [INAUDIBLE]. I think so, I'm not terribly sure. And in those days, that was more or less it. Dave Gross was around the company but he was doing writing some place else.

BL: What year was this?

AK: This was in 1966. And Tom Edgars was around the company but he was in field service. I don't remember the names of the technicians.

BL: Where were you located?

AK: We were on ____ ____ in the mill. In those days _____ _____ office was the very corner facing the mill pond and it was cold in there. All of 10 engineering was in the first sixth of that side of the building I think. You take the whole length of the building and our 10 engineering never did get larger than the first half of the building on that side and across the hall was Linc Engineering which later became LDP and the rest of the floor was manufacturing, some of it ____ manufacturing, some of it _____ manufacturing, a little bit of it, I think is even _____ manufacturing.

BL: The entire 10 group was 80 people?

AK: The entire 10 hardware group, in those days, was, - are we now 30 people including secretaries and technicians and the 10 software group which in those days were down with all the other sofware people on 12-2. If you can believe the entire software department lived on 12-2 and had the stockroom in the middle of the floor.

All the software people in the company fit on one floor of Building 12 and Larry Cortner(?) was the software manager in those days. [INAUDIBLE].

BL: How close did you work with the software people then?

AK: At that time, not at all. Now, not too much later than that, when they were still manufacturing KA10s at the other end of the floor with Bill _____. The 10 software people moved up across the hall and by that time, I was really engrossed in Kotok doing the processor logic for the KI _____. And we had a fair amount of interaction with software people. The software group by that time had grown to, I guess 30 people. When we were doing the KI _____, we wanted to calculate the delays on the various wires and the wire list and the wire list, among other things, would tell you where the wires were and coordinate system, so we could calculate because we knew what the coordinate system really meant we could calculate how long the wires were. We said, - well we obviously can do this with Fortran, well it turned out that Fortran had a bug, in those days, so that if you tried reading the wire list which looks very much like a huge deck of cards, it was actually a computer file but it was very strongly formatted like a deck of cards because that's where it came from. Fortran, T format

character didn't work and we tried backing it up and it got very confused. Kotok said, - well, let's try this new language that's just been put on the 10, - Cobal, this is the first cobal of the company, it was written by Al Blackenton. So we picked up the Cobal manual, started reading it and said, - well this looks pretty easy. Just read this stuff in and do this and do that. Sure enough, it seemed like it worked except that when we did compute this column times that column divided by that column, it didn't work. That's curious. It says right here in the book it's supposed to work. So we went across the hall, or actually the downstairs because Al B. worked in Accounting in those days, and said, - Al B., it doesn't work and we did compute this and that, just like it says in the manual, doesn't give us reasonable answers. He said, - you used the compute _____? Nobody who write Cobal programs uses the compute _____, _____ _____ never tested it. Thanks for testing it.

BL: So you started using ____, that was an interesting tool.

AK: I believe that tool probably was still in the [INAUDIBLE] ... It may have been superseded by something ___ ___ _____. Here, we were actually doing analysis of the wiring ... well, throw it in there and see what happens. Part of the problem was, the 10 line always

wound up with large back planes. The back plane for the KI10, and the KA10. The KA10 and the KI10 were both larger than the wire wrap machine could wrap in one pass. So, what you had to do in order to wrap the thing is you had to wrap the lower two-thirds of the panal and you had to wrap the middle two-thirds of the panal and you had to wrap the upper two-thirds of the panal and the restriction was that you couldn't have a wire that went from the very top to the very bottom because the wire wrap machine couldn't wrap it. And, of course, in those days the biggest panal that anybody else was wrapping was probably the PDP9, PDP15 panals which were still good sized but they were small enough to fit in the wire wrap machine on one pass. So we had to put a whole bunch of things in the wire wrap program that knew about this and warned you if you tried doing things like that and of course the PDP8 and later PDP11 stuff was a row of modules that were maybe 12 inches, maybe 15 inches wide and only one row only 10 inches high or something. So they never ran into ____ problem.

[END OF SIDE A].

AK: ... we were always _____ signal processing or wire wrapping. One of the things that was done for the KI10 was the _____ _____, semi-automated little old lady, as they called it. But doing wire wrap by hand was

incredibly error prone. When you're talking about the green _____, 36 pins on that much space with 30 gauge wires and all that, it was essentially impossible to wire 20 pins and get all of them right because you were sort of counting a _____. So Tom Stockybrand built this fixture that had a fixed aiming point for the wire wrapped gun and moved the back plane around so that the right panal was in front of the aiming point and ran off of a paper tape, I think later deck tape. So somebody could sit there with a, also had a read-out that said, - next wire you want is this length, so there were different lengths of wire already stripped, so you'd pick the wire out of the bin and wrap it wherever the machine said and touched the gun, ... little points of _____ at that end, the machine would move to wherever the other end was supposed to be and you'd wrap it so you didn't have to count all the silly pins. That was originally built for prototypes. But in the KI10, we decided we needed to wrap this pair, in order to have adequate signal integrity. Nobody had ever invented a machine that wrapped _____. [UNINTELLIGIBLE]. So Tom fixed the machine and we fixed the wire output program so that it would take, first the signal wire and then the ground wire at one end and twist it there and then the signal wire and the ground wire at the other end of the _____. So there were several hundred, maybe a thousand twisted pair in the back of a KI10 which could not have been

known without that fixturing tool. With the fixturing tool, it wasn't zero but it was [INAUDIBLE].

BL: Who was responsible for coming up with the design in the first place?

AK: Well the PDP10 was really a very good ____ of the PDP6. My understanding, and I wasn't here at the time so it's subject to correction, is that Gordon Bell and Alan Kotok came up with the original architecture. Now, by the time I was here, Gordon Bell was still here but by the time I was here, he had essentially nothing to do with any of it and Kotok was sort of the principal architect on the KA10, I guess Bob Clemmons did a fair amount of design with him. It seems like there must have been somebody else but I don't know who it was. By the time the KI10 came around, Bob Clemmons had left _____. So Kotok and Gross did most of the _____ part of it and I did a lot of the INO and we all did a lot of the [INAUDIBLE].

BL: Why was that?

AK: Because figuring out what he wanted to do was difficult and figuring out how to do it was even more difficult. There's just so many cases of _____. The KI10 was the first machine DEC ever built that really had

_____ in it. The 10 family has a lot of transfer instruction,which is willing to transfer, an almost arbitrary amount of data [INAUDIBLE] ... this address or that address. And the number of cases you get when the source of the destination reaches a _____ _____ and the next _____ is there, that kind of stuff will drive you up a wall. And we did eventually figure it all out. We designed it once and then built it and went off to design something else and then come back to it and said, - that's an ugly design and threw it out, designed it again. I think it was the third pass before we got something that we thought was good enough, that was worth building. Once we got it built, we only had a couple of bugs. That was a real pained(?) piece of logic. None of us were mathematicians and as a result we [INAUDIBLE] ... and the K__10 was off by a couple of bits, a couple of insignificant bits a fair amount of the time. Now, interestingly, the customers never noticed this. [INAUDIBLE]. When we designed the KI10 floating point a woman Mary Payne(?) had joined the company and she was a real dyed in the wool, honest to God, mathematician of the classical sort. She actually knew how to manipulate bits and stuff. She's still with us by the way. She said, - oh, well, here's how you really should do _____ point. And we said, - hmmm, those answers are different than what we're getting on the ____10. That's not good. Went back and looked at the KA10 and said, you know,

that thing is wrong. And we moaned and groaned for several months about whether we should fix it or not and then it was decided that, well, if nobody has complained, it ain't worth fixing. And this isn't as bad as it sounds because floating point arithmetic is, by it's nature, approximate. It was just a little more approximate than we thought it was.

BL: What project did you take on after KI10?

AK: Well, let's see, after the KI10, the KI10, that's about the time we moved from, afater we did the KI10, I started building communications stuff with Tom Edgars, a lot of different things. Tom designed the EX10, no Tom didn't design the EX10, somebody else designed it. Tom designed a, - the BL10(?) a connector, a rather complicated connector that maps four PDP11 memory spaces into the 10 memory and serves as a communications front end for the KI10 and KL10? The KL10 had an external memory. After Tom designed that, he had to make some change from 18 bit address into 22 bit address because 18 bits, which when the PDP6 was designed, seemed like all you'd ever need is only 256K worth of address. By that time, 256K worth of address was actually a little on the short side. So we added four more address bits in the KI10 fairly early. I guess the KI10 [INAUDIBLE]. ... which was a DEC ____ ____ to _____ channel bus converter

which was intended to be used on the 11 and the 10 [INAUDIBLE]. I don't know if we ever actually sold any on the 11s. They were eventually supported on the VAX ___. We just officially stopped supporting them a few months ago. They were the way in which you got a real pay drive, one that worked well, on their own or a 10 system. None of the tape drives DEC ever designed were brought in, it was really very good. They were as good as anybody else's in the industry except IBM but that wasn't very good. And when we started going in for a real business in our operations the customers said, - you got to be kidding. So we went out to Storage Technology Corporation, it was really a bunch of, IBM _____ who left IBM and _____ 71, 72 tape drives which ran on IBM buses and they were about to change them and provide [INAUDIBLE]. So we built this thing that sat on a _____ bus and from the other side looked like it was an IBM channel bus and to this you could then hook [INAUDIBLE] .. which were also built by STP and get real tape drives on it, a thousand dollars a drive onto your DEC System 10. The name had changed by then to DEC System 10. That was a _____ _____. About the same time, 10AM20(?) which was a memory interface. It was a memory interface to the Arbornet ___, got you directly on to the Arbornet. Both of those things were a little curious in that they were wire wrapped woods(?) because the volume was going to be pre_____, which it was and it wasn't worth the

money to build that _____, which at the time cost a small fortune. With a modern tool you could just do it anyway. With modern tools it's easier to do the X4(?). So we built those two things. We sold a fair number. I don't think we sold terribly many Arbornets.

BL: Did that also work with ANF10s?

AK: ANF10 was a set of software. That's another curious story. Dave McClure and the guy who is held up as the architect at that time, - Steve Wecker - sat down and designed DDCMP(?) designed for a company network where Wecker was working on the 11 space and the other guy was working on the 10 space. Then they went off and implemented it. Now, believe me, [INAUDIBLE]. And when the implementations were brought together because each of the impelementor, basically, when they found the problem, ____ cover the spec, fixed it. The implementations didn't play together at all. So DECnet, ____MP was the 11 version and what came to be called,- ANF10 was the 10 version and they both existed for several years side by side but weren't compatible. Eventually there was a 10 project to build a now compatible DECNET for _____ which is what's _____ for the last ten years. ANF10 was sort of the predecessor of DECNET and was originally intended to be compatible but turned out not to be. [INAUDIBLE]. After the ANF10 DECNET's with Steve Wilson(?) and I with

some assistance from Tony Locke and other people, sat down and tried to write a definitive specification for each DDCMP, the original DDCMP _____ was, as was obvious by then, subject to a number of interpretations. So Steve Wilson and I wrote a long, very definitive DDCMP _____. The people who were officially in charge of network said, - it's too long, it will scare everybody away. So it never, never got adopted or published but as recently as four or five years ago I ran into people who said, - who had a copy of it say, - that's wonderful. it's the only spec I've ever, ever seen that's really definitive as to what needs to be done. [INAUDIBLE]. It was the great type set 10.

BL: Type set 10?

AK: Yep, type set 10. By the time we did type set 10, type set 8 had been very successful. [INAUDIBLE]. I don't remember who got the idea but we decided that we were gong to do typesetting for big newspapers, big daily newspapers with the kind of _____ they needed. If you're putting to bed a daily newspaper and your computer goes down, you can't do any type setting, you're going to get really upset. So [INAUDIBLE] .... a number of other people involved writing the software. I the assistant integrator or engineer for reasons having nothing to do with DEC my dad, my grandad and his grandad were

printers. I knew a moderate to fair amount about type setting and printing in general and I could talk to these guys whereas the programmers said - we don't know. What's the point? Where's the product line? We put together a type set _____, put some hardware configurations and we got some _____ with TSS and all this and we actually sold three copies of it to the daily newspaper in West Palm Beach, the Kansas City Star, the Ontario Free Press. It turned out that no two of them were the same or even close, which meant that there was a lot of special work to go into customizing, and customizing the hardware. After those three installations we decided that was not a market we wanted to be in and we sold that business to some company who is really into typesetting. [INAUDIBLE]. The same software without all the reliability hardware was soon to be sold to the publishing system [INAUDIBLE]. The software really worked. The agony of having _____, for a small newspaper you can say, - well, this is the way you want to do it. And most of them will say, - yeah, okay, that's what I got to deal with. I ain't got a lot of money. Large newspapers say, - you got this great system and it's only a million and a half dollars, two million dollars, but I want it this way, that way, that way. That way, you can't sell it to me. You say, but, but, but, but. Why are you giving me this hard time about it? I mean my press unions cost five times that

much. Strange business. It turns out that the VAX, - there's at least one VAX system at the Boston Globe that's doing editorial composing and stuff.

BL: It seems that around that time, the seventies, you have different groups.

AK: Oh, no the 8s went on for a long time after that. People don't realize.

BL: [INAUDIBLE].

AK: The last that was called an 8, 8EAM, yeah. And then they turned into DECMATES and they lasted for another ten years or something.

BL: But there were four major engineering groups at this time; the 11 group, the 10 group, and the VAX group.

AK: It's more than that actually. There was the 8 engineering group, the 10 engineering group, the 7-9-15 engineering group, the 11 engineering group and the VAX engineering group and then there was a small, Russ Doane, was, I'm not sure, a lot of this happened before, apparently Russ Doane helped decide a PPDP6 module. By the time I came to DEC Rus Doane was the engineer on the 1000 Series module which were system modules that were

really [INAUDIBLE]. ... the internal products had set memory tasks(?). Then, somewhat later, he had gone to the other extreme with doing the ___ series modules, which are really, really slow so that they, in a noisy environment like that, not your controlled environment, they don't pick up _____ noise and go twitching around. What noise? I don't see any noise. So he had yet another group, so after the K series modules or as part of the K series modules, there was an industrial _____ group which was an interesting group but wasn't all that successful commercially and I believe he was _____. I can't keep the PDP14, PDP16 straight but I don't believe they were related. He was basically one of them and there was another one. So it was the entire market productwise(?) [INAUDIBLE] which was very, very successful. There were a lot of people who were buying modules to build their own ____. That didn't really die out until _____. Ken tells a different story or Ken's story has changed over the years and they're not really consistent. Ken's original story when I first joined the company was, - well, we started out building these modules and somebody realized one day that you can build a computer out of them so he built this computer. In his later story, which I have a feeling is closer to the truth, - well we started this little company, we really wanted to build a computer all along but we were smart enough to realize that we couldn't find the capital or

the engineering talent or the wherewithal to build a computer straight out of the box so we started building modules that developed the business. I'm not sure how you'd test between one and the other. They're both good stories. There's probably some truth to _____.

BL: I'm interested in all the dynamics of the different groups. For example, ocassionally you have to work together, as you mentioned. In what ways did the groups work together, in what ways were they closed and competitive against each other, trying to outdo each other?

AK: The markets tended to be very separate. I think the type setting market, the kind of customer you sold types of 8 systems to, obviously couldn't afford a 10 system and wouldn't know what to do with it anyway and vice versa. There were times in later days when the PDP11s got fairly large. It wasn't always as clear-cut as it might have been whether you'd be selling an 11 or 10 and either were facing a deal where we'd have friends that were now customers and say, - I just had your PDP11 salesman sell me an 11 and the day after I had your 10 salesman trying to tell me that the 11 salesman was all wet. Sell me a 10. There were problems like that. The other problem was that in those days we had product lines and each product line really did have its own salesforce,

it's own manufacturing allocations, in some case, it's own manufacturing line and some of the customers were smart enough to go product line shopping. They knew what they wanted. Which product line an give me the best valid[illegible]? _____ system, they couldn't do that. So the customer would say, - hey, you don't want my business? And of course that's very hard to turn down.

BL: In the engineering level, what interaction was there?

AK: Oh, the engineering, there wasn't an awful lot of interaction just because there was no need for it. A lot of the modules that were designed for the KA10 found themselves later [INAUDIBLE]. We all had the common catalog. A lot of the controllers for the PDP10 were built out of 8 modules. This is straight 8 modules. The PDP8 was built out of a whole bunch of modules in one board. The PDP1120(?) the first 11 was built out of a few special modules and a few 10 modules. And of course the next generation after that, all the smaller machines were built as a single module or two modules [INAUDIBLE]. At that point the commonality was, well, we're using the same ____ as you got. There wasn't a lot to be common any more. The wire wrapped tools and stuff like that we did with the 10 were generally used by a 1511(?). The 8 people didn't use [INAUDIBLE] .. they were irrelevant.

[INAUDIBLE]. The 9s, 10s, 13s(?) and early 11s were cabinet(?) machines. Since we weren't about to redesign the logic every time they need a cabinet _____. There was this mod(?) kit for the old cabinets that stretched it and had to put this.

[END OF SIDE B].

AK: So we got them to make to make this little excrusion that I don't if it screwed on or whether it taped on to the top of an old style cabinet to make it look like it was a file cabinet. This worked fine for production but some customers complained that they they had these old peripherals laying around from ____ 10 days and they got this new CPU and it didn't look right. So, we arranged to ship out to the customers who complained, just the mod kits so they could make their old cabinets look new. Then there was the great KI10 console flap. The industrial designers, presuming that the ____ of marketing decided that the KI10 would have a separate console. So they built this big chunk of sheet metal. It looked good and cost a lot of money presumably. They had this really strange idea for the casing, some sort of swirled yellow-orange paint scheme, which I suppose was very artisitc or something.

BL: That wasn't Corita designed one?

AK: It might have been. It was about that same era when we were doing Corita sign panals for _____. So we had this whole thing [INAUDIBLE]. And Ken came down to look at it. Ken sometimes did call himself the chief consul designer, probably after fiascos like this. And he took one look at it and said, - no way. And all the engineers said, - thank God. Put the panal back in the cabinet

where it should have been all along in the first place. So my love of industrial designers is limited.

BL: How did things change as it shifted into central engineering organization?

AK: Well, it's not clear it ever did. There's two issues: one is, until probably the 1980s, the PDP10 DEC System, hardware organization wasn't, quote, part of central engineering, unquote. It was part of the 10 product line. Now on paper, that's a big difference. In reality it's not ____ any difference at all. Likewise, there were all these, quote, central engineering groups. I guess by that time there was this fairly significant [INAUDIBLE]. There were people who would crow about central engineering being good or bad or wonderful, whatever, and the reality of day to day engineering is it didn't make any difference. Except from a manager's point of view it didn't matter whether it only worked for central (?) engineering or product line engineering or anything else. The engineers all talk to each other and there wasn't ... the only thing that caused engineers not to talk to each other was that it's about the same time that engineering started moving in the different physical buildings. And there wasn't yet any easynet or anything. So communications became sparse just because you could no longer walk down the hall and find somebody in their

office or meet them in cafeteria or whatever. that had a much larger effect than any central engineering product line engineering split. When I joined DEC, all engineering, in fact all of DEC except for the sales offices was in the Mill. In fact, it wasn't even all the Mill. Westminster was the first US separate plant, followed fairly soon by Salem, if I remember. But engineering didn't begin to split up until 10 (?) engineering moved to Marlboro in 1970, '71, '72, something like that. A year and a half after RISC (?) died. And I don't think any other engineering group moved until VAX engineering and VAX ____ engineering moved to Tewksbury. I think communications engineering moved about the same time.

BL: To the shopping center.

AK: Because the 11780 was designed in the Mill, and the 11750 was designed partly in the Mill and partly in Tewskbury.

BL: So you, as part of the 10 group, you went to Marlboro. Did that foster a sense of isolation then?

AK: To some extent.

BL: But you felt everyone else was isolated.

AK: To the extent that we ever had to talk to each other we still did. And in fact in those days there was a lot of traffic from the Mill to Marlboro and back. And in fact, I don't know what, I doubt if I can find it anymore. If you ever try getting from the Mill to Marlboro, you discover there's a large number of ways to get there and they all take 45 minutes. And at one time, Steve Ralston and I were going back and forth to Marlboro so often that I just started exploring all the routes. And somewhere there's a paper drawing to hand out to people that shows all the routes from Maynard to Marlboro, and with the comment of it doesn't matter which one you take, and it's physically true. But in terms of engineering it didn't seem to be any significant loss of communication because there wasn't an awful lot of communication between the groups needed. And the things that were needed to be common were common and the things that weren't, none of them cared. The 10 group, I trying to remember where the hell was CSS in those days. It must have been in the Mill somewhere, but I don't remember. But we had a number of strange system integration projects that I guess were too big for CSS or something, or we helped CSS with it or something strange. So one time we had a system in the Mill basement that consisted of one or two KA-10s. This must've been in 1968 or '9, and several PDP-15s and a bunch of PDP-8s.

And then as the PDP-11 came along, we basically -- I was the engineer on a lot of the remote station things, that are sort of like work stations of today except they didn't have tubes in them. Where we'd built PDP-8 and PDP-11 remote stations to go with your PDP-10. And unlike work stations, these were really remote. They might be several tens or even hundreds of miles away from the real machine. And this meant that instead of having a whole bunch of phone lines connecting your terminals to the machine, you'd have one phone line and a modem, and it would connect not only your terminals but paper tape reader and printer and I don't know if we ever _____ card punch or not. I don't think we ever officially put a card punch on one of the 8s or 11. And the software was there, but it was never officially supported. You could also have DEC tape on your 8 or 11. So in some very real sense we were doing distributive computing. Well, the computing was really all on the 10. But the users were not all in the same building as the 10. They were hundreds of miles away sometimes, long before work stations ever came around. And it seemed to me that when we first brought out VAXes we should have continued that and we didn't. Now at this point with work stations, it's all irrelevant to it. So even though we were in Marlboro we still had fair contacts with the 8 and 11 people, because we were doing their products and making modifications to their stuff.

BL: What became the engineering ethic at DEC? What was different about it in terms of management or structure, in terms of structuring jobs?

AK: Well, different from what? It's sort of the way I've always worked but I'm not at all sure that I've worked at the typical places. I'm not sure how to answer that. One of the things that perhaps distinguished the 10 group from the other engineering groups, was that -- apparently the management supported us but it was never all that clear. But the engineers felt that they were responsible for the entire system, making the customer happy. They weren't just a CPU group. And this was true until KL-10, but the KL-10 people basically took the tack of we're doing CPU's, we don't know about all this other stuff. That's overstating things a bit, but it's basically true. But the KA-10, the KI-10, the engineers felt responsible to be system engineers, which always makes me look a little funny at people who say, Well, DEC ought to have systems engineers. Well, we used to. And in fact the reason that the KL-10 didn't have system engineers was because management wasn't interested in doing system engineering, which was a mistake. Because what that really means is well if you don't do it, the field service people have to, the field service people, salespeople. And you're probably better equipped to do

it than they are. On the other hand, when you're selling the kinds of volumes we are today, you can't afford to have engineering doing it, you've got to push that out into the field.

BL: So everyone was a generalist.

AK: Yeah, I believe that's true. There were people whose primary responsibility was worrying about circuits. But they didn't have blinders on. I started out doing circuits, then did CPUs and peripherals and was doing systems along the way. There was no system engineer, there was Potok and me and Rose and Eggers, and all of us were doing system engineering as needed. And we'd all talk to customers as needed.

BL: In the '70s, DEC had a reputation of overengineering products.

AK: I guess that's true. I've never heard a customer complain about overengineering. Now maybe that's part of being too close to the products, but every product I've ever had anything to do with, I knew where the bugs were. And there always were some. I've never yet seen a perfect product. And all you can hope is you can get the bugs well enough hidden so the customers don't notice them. I've been reading some of the Six-Sigma (?) stuff

lately, and I've noticed that somebody seems to think that Six-Sigma is less than three lines of code in a million error or something. Or what's worse is I read somebody say that Six-Sigma is only three parts in a million on a functional circuit board. [UNCLEAR] That isn't high enough. What the customers want is no errors, no broken boards, whatever. So in that sense, I say, if that's what you mean by Six-Sigma that ain't good enough at all. Now if you mean by Six-Sigma there's only three broken boards in a million boards built, eh, that's better.

BL: Well, the 10s were pretty reliable.

AK: The 10s were pretty reliable. The Six was a disaster. The PDP-1 used a funny logic family that was raised on germanium. And used stacked transistors. A saturated germanium transistor only has about two-tenths of a _____ voltage. And so you could away with stacking three germanium transistors and still get reasonable logic _____. And the first PDP-6, serial number 1, was built with germanium transistors. Unsolveable (?) in my time. And somebody noticed that the PDP-6 generated enough heat that you really couldn't get away with doing germanium transistors. And so they switched to silicon transistors without changing the design noticeably. And this had the effect of making circuitry considerably

less, I don't want to say less reliable. It wasn't unreliable in the sense that it broke. It was unreliable in the sense that it didn't always give the right answer. And also the PDP-6 used a module of all the registers. So the basic PDP-6 was, the arithmetic part of the PDP-6 was 36 modules, and then control modules around, and stuff _____. That module was never testable by anything other than a PDP-6. There was a tester for it, and a tester would frequently say, this module's just fine, and you'd plug it in a PDP-6 and it didn't work. So they eventually threw the tester away. I don't think there's anybody who had anything to do with the PDP-6 who wouldn't say, well, yeah, that was a marginal design. In fact, late in JH-10 (?) they offered a trade-in deal to customers to get the PDP-6s back, because we figured we were better off not having them in the field, than having them in the field. There were only 22 of them built actually. There's always the Giant _____ story.

BL: What's the Giant Bryant (?) story?

AK: Well, back in '68 or '69, marketing decided we needed a big disk. This was basically true because the disk we had before was a Burroughs disk that was bought out as a compromise between PDP-9s and PDP-10s, and it was too expensive for the 9's and too small for the 10's. And somebody made this marvelous decision, so we needed a

bigger disk. And Bryant was _____, it was physically a huge disk. It had a horizontal shift, full capacity head, 13 or 14 platters, the platters were two feet in diameter, maybe larger. And that was the mechanical assembly. Then there was an electronic assembly that sat next to it that had cables over to the heads. This whole thing was hydraulically actuated. There was a hydraulic power supply to supply hydraulic power to the hydraulic action, and then there was an optional cooling. Oh, the thing was water cooled. So there was an optional cooling unit that supplied chilled water to it, which if you already had chilled water in your building you didn't need, but almost everybody did. And we did all the engineering, we built a controller for the damn thing. It was the RP-10. The disk was the RA-10, the controller was the RP-10, I think. So we did all the engineering, we had it working in the lab. We shipped it out to the first customer, the Computer Center Group. It arrived two days after they went Chapter 11, so they never paid us for it. And we never managed to sell another one. Because at the same time IBM had brought out the 2311 disk-pack, and everybody was saying, that's what you ought to have, sell me one of those. That's what we had been telling marketing all along. And they said, no, no, no, we want this giant fixed disk. Their reasoning wasn't entirely unreasonable, because Applied Logic in Princeton, New Jersey, had three or four

sixes that they were selling time sharing service on, and they had gone to Bryant and they had done the engineering and they had done this disk on a PDP-6.

BL: This is Brian?

AK: Bryant. B-R-Y-A-N-T. _____, Minnesota. They may still be there for all I know. I'm sure they're not selling this anymore. And Applied Logic had actually had good success with this disk. But of course, they had started using a 3 years before or something. And by this time IBM had brought out the 2311 disk-pack, which was small, reasonably convenient, had interchangeable disks, and engineering saw what was happening, and said that's clearly what we should do, we should build an interface for that thing and not raise Giant Bryant, and marketing says, no, no, no. So as it turns out, we were right, marketing was wrong. Sold one, got paid for zero. And then we had this crash project to build a controller for the 2311, 2314, which became the RP-10 and the RPO-1,2,3.

BL: What happened then. Were you involved in Jupiter design?

AK: No, by that time ... where was I [UNCLEAR]. About 1970, I don't know, 1 or 2, after we did the EM-10, and AM-20, I went over and started working on _____,

originally trying to get engineering more organized in card numbers and stuff like that, and eventually wound up working on Eplus (?) for several years. Eplus is the DEC system 10 predecessor of Apix which is on the VAX. And so I was sort of the interface between engineering and drafting type users and the programmers and trying to help the programmers figure out what the right thing to do was, and trying to help the engineers figure out how to talk to the programmers and whatnot. During the same year I was running engineering, so I was not out of touch with resident engineering.

BL: What was it like running the engineering committee?

AK: I'm not sure how to answer that. Before I took over the engineering committee, in fact before I was even really associated, the engineering committee basically passed on all the engineering projets and it was Ken's engineering committee, quite literally. And by the time that I got involved with it, the engineering project reporting and budgeting and stuff had very appropriately gone off to the engineering manager who was responsible for it. And the engineering committee was basically trying to fill the cracks between the engineering groups and between engineering groups and support groups and engineering and manufacturing and engineering and field service. In the process, well, the, being responsible

for the entire VAX standard system. Some people say that's all they did, well no. Standards were only a means to the end and the end was what was important. Because a lot of people probably don't to this day understand that. But basically, somebody would bring up a problem and we'd try to get the people in there, and try to get the problem solved. And the problem might be something like one engineering group is doing something that's different than another engineering group, but they all have to be read by the same field service people or ... Some equipment was probably still in the field and some was .... I don't know what else. One day Ken came down and said, I want you to get the DEC standard 200 people moving. Well, we tried. And to this day, the DEC standard 200 is the whole communications standard, and I think to this day they haven't been revised in five years or something. Oh well, you can't win them all. The engineering committee basically dug into anything that somebody said, hey that's a problem and it seems to be related to engineering. I don't know the company's probably outgrown the function that needs to be done and being done in ad hoc fashion in various places then.

BL: Again Ken (?) it seems make anybody or any group do anything in Digital. So you formed a negotiative role?

AK: Yeah, in the cleanest way of our role of the

engineering committee's role was to get the people who thought somebody wasn't doing something for them, and the people who were supposed to be doing it, and try to get them to sit down in one place and then help them negotiate. And frequently it wasn't a two way deal, it was a three or four way deal or something. It probably worked a lot better when everybody was in greater Maynard. This was before Colorado Springs, or before there was major engineering in Europe or major engineering in California and Oregon. And it's also before there was Easynet. And to a large extent, Easynet has made a lot of that negotiation simpler and easier for a lot of people to participate in. And sort of the march of technology has made a lot of things come out for free, that is to say, this is the only design tool that really does the job, so everybody's sort of stuck with using it, whether you want to or not. I don't have to tell you you've got to use it, but there aren't any alternatives that we're talking about. And a lot of it is automation sort of forces you to do things the same rather than be gratuitously different. Like there's the whole module design process. In the days when you were laying out modules by putting pieces of tape on mylar and manufacturing would photograph them, well, it didn't really matter what kind of design tool you used, because manufacturing never saw them. It became true, good God, certainly by 1973 or '4, that pieces of tape on paper or

mylar didn't have enough resolution to build modules with anymore. So we had to do modules using some CAD system. Well, we've gone through several CAD systems over the years, but basically only one at a time, so whatever the CAD system did was what you were stuck with. And I've never yet seen a CAD system that did everything that everybody wanted. So it was always a matter of being stuck with it. But that means that manufacturing can build an interface of the CAD system output and whatever was there was what they got and they couldn't get anymore and so they had to live with them a lot, but they knew what they were going to get, get the very groups together. So automation helps standardization an awful lot. Because people mumble and grumble, but if the choice is doing it by hand or doing it using a tool, the tool usually wins out. And eventually manufacturing comes over and says, well, I'm not going to take it by hand anymore anyway.

BL: Where did the DEC Standards come from?

AK: The original DEC standards, which have been lost actually -- there's an old DEC Standard 2 and an old DEC Standard 3, I think I even had copies at one time, but I lost them, were done by Harlan Anderson and somebody else. Long before I ever got to DEC. [UNCLEAR] I'm not sure what the original one was. The current one is how

the DEC standards system works, which is sort of irrational (?) but I haven't _____, that wasn't the original DEC Standard 1, because originally there was one. And I had no idea -- _____ might remember. Slowly but surely as we put things on computer, we'd get in control. I had a call from the standards people a month ago maybe, saying, We've got somebody who wants to use the DEC Standard 11 number for something, and we can't find any trace of that standard ever having been used before. And I said, well, wait a minute, and I pulled out my notebook and said, DEC standard 11 is how to draw something or other symbols and I've got a copy of it right here.

[END OF SIDE 3]

AK: It turns out the DEC Standard 11 is obsolete and has been for probably ten years for the design group, and three or four years for service purposes. But there is still hardware out in the field documented using the DEC Standard 11 symbology, so you don't really want to throw it away and say due to our wonderful ability to keep papers over the changes of personnel and groups and whatnot, the Standards and message people didn't even have a copy of it. It turns out the copy I have wasn't even complete. But they didn't even have any record of it. That happened. Another neat story which has nothing

to do with anything I had to do with. I don't know if you know what aperture cards are, but all the drawings that go through drafting and get signed off, get microfilmed on 35 millimeter film. And the cell gets cut into the frames, and the frames are mounted on _____ compatible tab cards with a hole in them, which are called aperture cards. Now originals for these things are put on 35 millimeter rolls of microfilm, it's the same size as motion picture film. The spools are only about four inches in diameter. The archival copy of these was apparently stored in the one of the Maynard banks years ago, in the basement. But one day, drafting in Maynard got a call, or actually I suppose the window in Maynard which in those days did all the microfilming, got a call from the bank saying, somebody overfilled our oil tank and it didn't stop, so all your microfilm has been oil soaked. And I don't know whether they managed to recover that or not. But if they didn't, there's probably a lot of old data that got lost. Now, the other piece of archiving which as far as I know is still not fixed, is the people who run, who make the aperture cards and whatnot, distribute them around to various plants like Marlboro and Maynard and _____, I guess, I'm not sure where else. Where the so called co-files of just tons of these aperture cards. Well, Marlboro for example sees an aperture card and says, anybody interested in this one? No. Toss. (?) And the people who were

running the show never realized they did this. So they were saying, even if one of these files goes up in smoke or somebody bombs the building, or a fire comes in, it doesn't matter because we've got these other files all over the world. Well, the other files are probably not nearly as coherent as you might hope. And a lot of these files are at manufacturing plants, and I'm sure the manufacturing plants take the attitude of, well, we're not manufacturing this anybody, goodbye. And of course we service things for sometimes 20 years after we manufacture them. So we could be in trouble on some older products.

BL: I'm sure they'll never fail. They'll keep going forever.

AK: So far, I've never heard of us actually getting in trouble, but we've been loose with keeping our archives over the years. Always a little troublesome now. The modern version of that is well it's all in the computer tape right? We started the computer tape in Iron Mountain (?) right? Yeah, right. Is anybody worried about rewinding the computer tape and making sure it's still readable every now and again? I'm sure.

BL: So the last few years you've been doing systems integration?

AK: Well, it's sort of curious. The catch up beyond Eplus, after Eplus, actually during Eplus, I was talking to a guy who I guess is still around by the name of _____, who was a product manager of the 10 world in those days. And he had come from Honeywell where he had a configuration (?) tool. So I wrote -- apparently the third experiment on the configuration tool at DEC, as I learned later, were a couple of other attempts earlier to either I wrote a configuration tool for the KL-10 machines, and it was actually in limited use in the field for six or eight months, but I made it quite clear this was a prototype and it was really -- now that I demonstrated it was possible, somebody ought to sit down and write it for real, in a maintainable fashion, because it clearly wasn't maintainable and I never claimed it was. And nobody did that, so that fell into disuse. And in the meantime I had joined Arnold Kraftman at Sigma in the beginnings of what's now the XCOM (?) project. And worked not as a programmer in that case but again as an engineering, as an interface between the programmers and the engineers who knew the products. And we got XCOM and XCEL probably going. [UNCLEAR] And then I worked in what was called the International Product Requirements office or several years doing, trying to get our act together on what it was needed to build products that really met international requirements. For years we had

been shipping products to Europe for instance with American power cords. And instead of getting the right power cords, we merely painted them red to say this is really a 220 volt power cord, not a 110 volt power cord like it looks like. Well, eventually customers decided this was sort of ridiculous, I helped with Frank _____ and some other people to bring in all the power plugs that are used around the world so that we now ship terminals and small systems and whatever with the right power cord for whatever country it's going to. Which turns out to be a problem for some countries like Saudi Arabia where which power cord you need depends on who built the building it's going in. If it was built by the Germans it's a German power cord, if it was built by an American it was an American power cord. So I worked there for three years and then for another three or four years I worked over back in the old 10 space, except by now it was the VAX Clusters space, originally doing customer support and planning and then being the architect on the CI-Switch (?) product which we spent three years on and then decided we didn't need after all. In retrospect I think the decision was right, but it's sort of a nuisance of building something that large and getting to the point of actually having working prototypes and then scrapping the project.

BL: Are many products developed that never see the light

of day?

AK: It's hard to know. There certainly have been a lot of products that are proposed that never get any further than that. The number of products that are actually brought to hardware and then don't get shipped is probably not very many. Probably less than five percent, I would think. The big one was Jupiter of course, which I wasn't in any way involved in. I was working in Marlboro those days, but I knew people who were. That was a case where it didn't work. And they never got to the point of actually building hardware, but they couldn't figure out how to design it so it would work so they gave up on it. Unfortunately this was after telling about a hundred customers that they were going to get it. That's an awful lot of customers. In fact the same thing was true with CI-Switch although not quite to the same extent. And we had gone out and talked to in one way or another maybe as many as a hundred customers, and some of them were enthusiastic about it, and some of them we set up arrangements with, and then cancelled it. A few of them, well, most of them thought they really needed it, and a few of them maybe genuinely needed it, and a lot of them, as it turns out -- When we started building the CI Switch, we thought _____ the bandwidth and perhaps more interconnectivity. The feature that was really the most appealing to customers was the ability to

take different things connected with the CI Switch and have them on logically isolated CIs. So you can take a CPU that was on one cluster in the morning, and say well, I don't need it there, I need it in this other cluster this afternoon and switch it. This was an idea that actually came up by Steve _____, who was a product manager type person. And he said, could you do this? I think we can figure out how to do that, that's a neat idea. That was what really the most saleable feature of the thing intitially. Now, when the 9000 came along, we knew the 9000 was coming along, we'd say, with the 9000 you're at least in bandwidth too. But when you were talking to customer, bandwidth? We don't need bandwidth. We've got enough bandwidth for all we're willing to pay, but that switching ability is really neat. And there are a few customers who have had to come up with various kludges to do the switching on their own, because we haven't had any product to do that. And it turns out the bandwidth problem doesn't seem to be much of a problem. With the 9000 what we finally did was make some minor changes in the VMS so we'd be able to support multiple CIs. But I'm not that close to the 9000 shipments or anything, but I believe the most we've every shipped on a 9000 _____ CIs. And that's usually more than enough. So the original analysis to say we needed 8, 10, maybe 20 times the bandwidth appears to have been an error. But the idea of being able to reconfigure your system without

shoving all the cables around was a real saleable feature that we don't have anything to substitute for. And I spent basically a year and a half, two years after the CI Switch, unselling the CI Switch, selling other documenting for the sales force, other alternatives to do it, and basically doing VAX Cluster configuration documentation. [UNCLEAR] almost exactly a year ago. And I haven't really gotten away from Clusters, I don't know if you know anything about system integration engineering, but we've got this large lab up in Salem, which is actually four logical labs in one big large space in Salem. And our group does _____ testing. [UNCLEAR] And we're doing a large cluster testing, originally for DuPont, and all of a sudden under arrangement with various other people _____, the VAX Cluster group in Marlboro. But we have the space and the facilities to do good stable testing that aren't generally available in other places. And we're currently negotiating with some of the Marlboro people that has a very large cluster, larger than anybody's ever done before. And I think we know where things break so we can know how to fix them until we try it.

BL: The world's largest VAX cluster (?).

AK: Oh yeah. On paper a cluster consisting of 16 VAX 9000s and hopefully without any changes at all, 224 total

units, most of which are work stations, should work.

BL: I remember the original spec was 16. And then it grew to 64.

AK: Well, actually it grew to 46, for reasons that are somewhat obscure. And then it went to 96. And in fact we have a couple of customers that are actually running 100 and 204. We believe there's probably some code in DMS that was perhaps a little shortsighted that will cause us to stop working at either 223 or 255. And we believe that even that code is probably _____.

[TAPE CUT?]

AK: -- that's what it was called. See, the PDP-6 when it was built back in 19 ... '65. Nobody had ever seen an operating system last as long as ten years. And in those days, memory was tight, and bits on tape were hard to come by, so the PDP-6 operating system had a date format that basically only used 12 bits and encoded things as the number of days in a funny kind of code, saying January 1, 19, five or something like that. Well, if you do a little calculations, you'll discover that 12 bits runs out at the end of late January or early April, something 1975. Well, I mean the Six operating system is really what turned into Top 10, in all the years, and

this date format was still embedded in the code. In fact it was embedded all over the place. It was on DEC tapes and _____ tapes, and hidden in this program and that program and it was all over the place. It was the standard way the operating system passed the date in applications and everything else. Somebody of course woke up a year and a half earlier and said, You know, this is going to run out and we're still going to be shipping this stuff, and we've got to do something about that. Well, most places it wasn't too bad. The rest of the field was blank, empty anyways so you just took the 12 bits and expanded it to 18 and that was plenty. Probably the touchiest situation was the DEC tape where there weren't more bits to expand, so you found an odd bit here and odd bit here and an odd bit here, and concatenated those to 15 bits, which was really enough. And we ran tests, and we set the clock in Marlboro up six months or something and we ran tests, and we'd find problems and we'd fix them. And when the date finally came, and we thought we had it all fixed, and we did pretty much, but there were three obscure products that were only used by a couple of customers that hadn't been fixed.

BL: Did you hear about it?

AK: Oh yeah, we heard about it real quick. The customer

would say, It doesn't run anymore, or it thinks it's 1965.

BL: I heard some other date stories, that the year 2000 will be trouble for some other numbering schemes.

AK: Well, it's not clear. The obvious place you'd expect to get into trouble, doesn't actually happen in 2000, and that's leap years. Because the way leap years are calculated is that it's divisible by 4, it's a leap year, unless it's divisible by 400, and if by 100, in which case it's not a leap year, unless it's divisible by 400 which it is. Well, 2000 is divisible by 400 so there's no break in the pattern. On the other hand, I suspect that all kinds of things are going to break, because they think that the year, by God it starts with 19, and it's not going to anymore.

[END OF INTERVIEW]