

An Interview with
ELDON WORLEY

OH 407

Conducted by Jeffrey R. Yost

on

15 November 2012

Computer Security History Project

Corvallis, Oregon

Charles Babbage Institute
Center for the History of Information Technology
University of Minnesota, Minneapolis
Copyright, Charles Babbage Institute

Eldon Worley Interview

15 November 2012

Oral History 407

Abstract

Eldon Worley was a longtime software scientist/engineer at IBM Research in San Jose. He pioneered Information Management Facility (IMF), a path breaking security system that was the underlying basis for the IBM security product Resource Access Control Facility, or RACF. After this work left IBM Research to be developed into a product, Worley continued to analyze and provide feedback to IBM's development staff. RACF has gone through many different releases since its first iteration in the mid-1970s. Worley provides background on the origins of IMF and RACF, and how RACF evolved over two decades—including perspective on customer installations of RACF in the UK, where he spent time in both the mid-1970s and mid-1980s.

This material is based upon work supported by the National Science Foundation under Grant No. 1116862, "Building an Infrastructure for Computer Security History."

Yost: My name is Jeffrey Yost, from the University of Minnesota. I'm here today in Corvallis, Oregon. It's November 15, 2012, and I'm here at the home of Eldon Worley. Eldon, can you begin by just telling me when and where you were born? And where you grew up?

Worley: Actually, I was born here in Corvallis in 1940. Grew up here, basically, until I graduated from OSU, so I'm a native, I'm a real green person. Married my wife in 1958 just a little bit after we graduated from high school. We knew each other in high school. Graduated as an electrical engineer, and at that time, Oregon State was anti-IBM, which kind of caused me not to interview IBM. So I went to work for Westinghouse and we went off to . . . (pause)

Yost: It was electrical engineering you studied?

Worley: Yes. Yes, although I was the maintenance engineer for the computer here at OSU, which was in the math department. I did that for three years or so, which was really good experience. It was a hexadecimal machine, which turned out to be useful.

Yost: And what type of work were you doing as a student?

Worley: Well, basically maintenance of the computer; when it didn't work, fix it; we added new instructions to it. The university owned the computer so we could kind of do anything to it we wanted. This was back in the days of the tube computer, if you can

remember such things. And we changed a lot of the gate logic to be transistorized. We built the boards kind of by hand.

Yost: This was an ALWAC?

Worley: Yes. It had no high level language. You had to know absolute assembler; there was no assembler so everything was done in hex, and paper tape input, or a typewriter. Didn't have a card reader. This was back in the really early days of the computer. It was difficult to make it continually running in the summer because it got hot and we didn't really have good cooling. It was sort of in the half basement of the math department so summers were always a challenge. We often had many fans in there to keep it cool. It was good experience, though.

Yost: And what led you to apply to Westinghouse after that?

Worley: In those days, you just sort of interviewed a bunch of companies and Westinghouse sounded pretty good and they had some work in computers. They were working on a parallel processing computer, which eventually ended up going to the University of Indiana, or Illinois, I'm not sure which. I was hired to work in the arithmetic unit design. We moved to the East coast and I spent three months in Pittsburgh, because they decided that they wanted me to go to a graduate school. I don't know how they picked people but I got picked so I spent three months in Pittsburgh and my wife came back and lived in Oregon because I was basically gone. And then I went

back to Westinghouse south of Baltimore and started working on the arithmetic unit, floating point arithmetic unit design. Then about 11 months from the total time with Westinghouse, Westinghouse decided they didn't want to do computers anymore. So I got moved over into a radar group, which I did not like. I just had a problem with working in that group, it wasn't at all what I wanted to do. And about that time there was a job fair down in Washington, D.C. and there were several of us that had been moved out of the computer group. So a bunch of us went down to Washington, D.C. and interviewed. I interviewed IBM and couple days later they said yes, come and see us. And so they flew me up to Endicott in New York, and I got a job there working in writing logic simulators for various pieces of hardware that Endicott makes.

Yost: This was in 1964?

Worley: This is, let's see, yes, the end of 1964, beginning of 1965. Yes.

Yost: Right about the time the System 360 was coming out?

Worley: Yes. I saw the first System 360s without even covers. There was a building down the street from where I worked that had a System 360 in it and, I don't know, somehow I convinced my manager that we ought to go down and participate in that. And so I worked there and got familiar with it. The operating system was really on cards and then there was DOS and there was TOS. In those days you had to be very careful when you walked into the computer room because of static electricity so there were these mats

on the floor and every once in a while somebody would come in and for whatever reason, touch something and pffft, the whole computer would go down. So I learned those basic operating systems. Then, for whatever reason, I got interested in understanding PL/1, because that was just at the beginning of PL/1, and so I became intimately familiar with how it worked, and how it expanded code, and that kind of thing. And at the same time I was doing a real job, regular job, of writing logic simulators.

Yost: So that was just an interest of yours?

Worley: Yes, but my management put up with it and eventually I kind of switched over into the 360; knowing the software, knowing JCL, which almost nobody knew at that time, and understanding PL/1, and kind of became a reasonably knowledgeable person in that area.; when I first went to work at Endicott I told my manager that probably in three or four years I would want to go back out to California, I would want to go back to the West coast. I went to a SHARE meeting in San Francisco. I don't know how I got involved in SHARE, but I did. And there was a friend I knew from the California lab, and I talked to him about maybe coming out there and he said well, let me talk to my people. In those days, IBM meant "I've Been Moved" and I got an offer that said hey, we'd like you to come out and you could be the person that teaches our people about System 360 stuff, and PL/1, and all of that kind of stuff, and JCL; and I said okay, I can do that. So they moved us out. We had just bought a house in Endicott, and 11 months later we had to sell it, but that was standard IBM operation. So we moved to San Jose, and at that time,

I worked in the manufacturing division. Basically, it was all software that I did; teaching software; teaching during the day; teaching at night (pause)

Yost: Was this both teaching other IBM employees, as well as the customer base?

Worley: No, strictly IBM. This was internal at this point and I guess a lot of what I said couldn't have been customer-oriented because I knew a lot of the internals. I knew how things expanded, and I knew what not to do, and what to do; and I did that for, I don't know, three, four years; and I became friends with what was the functional manager in the area I was in; he was about three levels above me. And somehow, he was friends with the IBM Research Division side, equivalent to him. And I don't know, one day I mentioned to him I'd really be interested in getting over into research. I guess it pays to know certain people. They arranged for me to go over there and work. In those days, you worked on a trial basis; you never officially transferred into research, until some period of time went by and they decided you were acceptable, whatever that means. So they gave me this project that came from Yorktown, called IMP, Installation Management Package.

Yost: That was in 1970 that you moved to IBM research?

Worley: Yes, right.

Yost: Prior to that, had you had any reason to even think about or consider computer security?

Worley: No. About the only thing that existed in those days was passwords and that was it.

Yost: Right. Well, MULTICS had some security designed in.

Worley: Yes. And I was familiar with MULTICS.

Yost: In fact that was programmed in PL/1.

Worley: Yes. Which people said was a bad thing, but I actually think was a good thing. It's just that PL/1 implementation wasn't very good. But the idea of IBM using a high level language I think was a good idea, much better than writing in a simpler language.

Yost: So you joined IBM Research Division in 1970.

Worley: Yes.

Yost: And that was Santa Teresa?

Worley: No, that was San Jose. At that point, they'd just built the building there; it was a triangle; it was shaped in a triangle. And that's where I joined, worked on this thing that we got from Yorktown. It was my task to get it to work, and I did in three or four months. My manager at that time said, we'd like to offer you a permanent position. So I became a permanent research member, which was a really good thing; best place in IBM to work. Maybe the best place in any; well, research is a great place to work, period. But it was kind of the start of RACF; or RAC, as it was originally called.

Yost: So there was an installation management project, and then that was expanded into what became RACF?

Worley: Yes. Yorktown had had it for I don't know for how long, but it had been a disaster there. It kept taking their system down. (Laughs.) They convinced San Jose Research that this is a great project, it works really well, and you won't have a problem with it, and it's yours. Later we found out that it was an absolute disaster and they were running a highly modified JES3 system and the two didn't get along well. It never really; well, after I altered it a little bit, it never really failed. We were using, at that time, it was a 195, and so I had to be really careful what I did. I couldn't take the machine down because there were jobs that would run several days and the people that ran these jobs were highly visible, right up to the president of IBM, and they thought nothing of calling up the president and saying you gotta get rid of this guy, he's keeping our machine down. So you were careful. (Laughs.) Or I was careful. People kind of came and went. I guess I was the only one who sort of stayed with it. There were maybe two or three other people.

One guy eventually quit. The other guy retired. Third person, I don't know where he went. People sort of drifted off, but I stayed because by the time we got it running, or I got it running — I always like to say “we”, it distributes the blame — by the time we got it running, some years had passed and the main manager in research, he said, you know we try to get things so that five years from when we start it'll be product level. I think it was about four that I made it; we got it to be considered as a corporate thing. We must have quadrupled in the lines of code from the original thing we got from Yorktown, because the original thing we got from Yorktown basically did hardly anything. So we wrote TSO commands; we were probably the first people ever to write TSO commands other than the famous ACCOUNT command that IBM wrote. There was really no design, no specs for how to write a TSO command so we kind of did that from trial and error. I remember stealing, um, borrowing the password algorithm from APL, because there was a Model 75 in the same computer room and we accessed the source. I was looking around saying, well, I need something to encrypt passwords and somebody said well, go look at the APL code. I looked at the APL code and that's why the first password encryption RACF did was so bad, because the APL thing could be reversed without much trouble. But that was all we had. Things like DES just didn't exist yet. That's kind of how it evolved.

Yost: Speaking of DES, IBM had by that time started a research program dealing with encryption. Was there any connection between that research and your work?

Worley: No. That was; there's an East coast/West coast culture in IBM. Still is, I think. And the East and the West coast often don't talk very much; in fact, probably in some cases, don't even like each other, which is just based on history and not reality. Once I got into the SHARE business, then I began to meet people from the East coast, the development people in security. Bill McPhee, Bob Courtney; met them at a SHARE meeting in Montreal, which is where I presented what I had worked on, and Barry Schragger presented what he had done, and some guy neither Barry nor I can remember who we think came from Boeing, but we don't know his name. We both tried very hard to figure out his name and we can't think of it.

Yost: How did that session come about?

Worley: My feeling; what I heard later — not at the time, but — at the time I was just invited to present a piece of a session in front of customers that discussed this “thing” that IBM had done, and it was considered a corporate project or a corporate product. Later I found out that the whole session was engineered by IBM development because they wanted to see how customers would respond to these three things and in particular, they wanted to see how they would respond to the IBM work. Also found out later that they; there was an effort in IBM, in Gaithersburg, working on something called RAC. But it really didn't exist; it was just a paper tiger; it was a spec about like that. Some work had been done out in the West coast in data management. They had interfaces that would invoke this service that didn't exist and there was, my understanding, at least one or two IBM vice presidents that weren't very happy with how long it was taking. And they felt

that IBM needed to be a presence in the security business. So that kind of led them to causing this presentation in Montreal. Shortly after that, I got invited to Poughkeepsie and several meetings with various high level managers. Mostly I was just a presence. I was just there to listen. I had been given instruction; just sit there; just act like a nice research person; don't do any; don't rock the boat or anything. I said okay, fine. I was in the meeting where this IBM vice president whose name I don't recall, was on the phone to the Gaithersburg people and basically read them the riot act about why they hadn't done anything, and how much money they had consumed with doing nothing. At the end of the call he finally said well, we're going with this thing from research; that's the way we're going to go. And I thought, wow, what?! (Laughs.) I mean, it's really nice to get a product in there but that's a lot easier said than done. Then I began to fly back and forth from the West coast to the East coast, and spent a lot of time talking to Bill McPhee, and Courtney, and at the same time trying to do the corporate thing back in San Jose.

Yost: Before that Toronto meeting and that session, I understand that Barry Schragger was asked to lead a task group.

Worley: Yes, and I didn't know anything about that.

Yost: Okay, so you hadn't been part of that task group.

Worley: No.

Yost: Did you meet him at the Toronto SHARE?

Worley: I met him at the SHARE meeting but that's the first time I'd met him.

Yost: What was your impression of ACF that he presented, compared with what you had been working on?

Worley: They were different. They were quite different. I remember him making comments about well, we use VSAM to store things and RACF did not. Although RACF used an organization that came from AM1, which was a VSAM predecessor, which I never said anything to him about but, you know. There were lots of reasons why we didn't use VSAM. That came up frequently from the data management people, "you ought to use VSAM." Well, I had a whole list of reasons why I couldn't. "Well, we're going to fix that." Yes, well the check's in the mail but I can't wait. I didn't know much about ACF2. In fact, I was directly informed not to look at any manuals, any publications, at all. We don't want any product contamination, which I understand for reasons of lawsuits and things like that. So I never really talked to Barry until after I left IBM and he and I both worked at Vanguard, and that's when I talked to him a lot more. To this day I know very little about ACF2 or Top Secret.

Yost: McPhee and Courtney were trying to gauge the reaction of customers?

Worley: Yes, are customers interested in this? Is it something we should pursue? And apparently (pause)

Yost: Did you get any sense of the impressions they had? And what impressions did you have about customer interest?

Worley: “They” meaning customers or . . . ?

Yost: Did you get a sense of McPhee and Courtney’s reading of the customers?

Worley: No.

Yost: Did you have a sense of what customers thought?

Worley: Customers asked a lot of questions. That, I thought, was interesting.

Yost: Do you recall any of those questions or the type of questions they asked?

Worley: I don’t, I’m sorry, that was just a long time ago.

Yost: Definitely was. You sort of answered this question, but were there any models or research out there that you found helpful in developing RACF or were you working in isolation?

Worley: Well, okay. I was constrained by the spec that Gaithersburg produced, somewhat. I had done a lot of looking at the Orange Books but the big criteria when I went to Poughkeepsie was 'lines of code'. At that time, that was a big deal in IBM. We can't have any extensive amount of addition of lines of code. They did some pruning of things; basically had to present to them all the things, what the use was, why they were put there, where maybe it might go. There were two guys in Poughkeepsie, we used to call them the Ewing twins. They weren't twins, they were brothers, and they had been associated with the Gaithersburg effort so they often asked questions from that perspective. And sometimes they were, well, it doesn't do this, that means it wasn't as good. I just didn't answer those; I didn't have to because there was McPhee and Courtney and other people there in the audience that dealt with that kind of stuff, and basically saying well, this is where we're going. The other one's dead.

Yost: Were you aware of a survey that IBM did in the early to mid-1970s on what customers; how much of their data they felt was important to protect?

Worley: Well, I became aware of it after we started working on the transfer of the code. There was the famous "Ten Percent Statement." There was a big controversy in the early development. Data management people, they wanted a bit on each data set descriptor. The bit said that there's protection. The original code had no bit. Everything was protected, period. And that was also consistent with what ACF2 did. But there was a lady in the development side who said well, there was this study or survey that said only 10

percent of customer data will ever be protected, so the bit isn't going to be a big deal. We're just not going to have the overhead problem. There was some loud discussions about that but she won, so the bit appeared and that part of the code got put back to being bit-driven, which looking back — hindsight's wonderful — was a terrible decision.

Yost: And this was to conserve memory resources, reduce overhead costs?

Worley: Yes, right. The idea was, well, if the bit's on, you go check the security. Bit's not on, there's no overhead at all. That was the reasoning. And because there's only 10 percent of the data protected, we're not going to get a lot of overhead. That wins in some manager's minds. I never saw that survey, though. It would've been interesting to see it but I never saw it. I just heard about it.

Yost: And did you also — you talked about McPhee and Courtney — did you also have any interaction with Bill Murray in this time period?

Worley: Yes. Yes.

Yost: In what capacity did you interact?

Worley: He was generally there in some of the meetings. There were three guys in the Montreal meeting and one of them I can't remember. He was a manager. There was McPhee, and Courtney, and this third guy. It was not Bill Murray; that I know. But Bill

Murray, yes, I talked to him several times after I ended up in Poughkeepsie on a temporary assignment. He's an interesting guy.

Yost: And so you had created a system with protection by default but that was removed.

Worley: That was removed, yes.

Yost: As I recall, there were four primary features, in the document that you sent to me, that were removed in creating the RACF product. I'd like to go through each of those and get your comment. Caching of profiles to speed up processing?

Worley: Yes. I didn't do that. A guy who was in ASDD at that time, a division that eventually went bye bye. He came into research, but; his name was John Palmer, and he did that because he felt that it was important to cache the stuff, that would cut some overhead down. I said yes, I agree; do it and I'll stick it in. So we had the caching and when you run in basically a protect-all mode, caching is important because you can go run through memory a lot quicker than you can do one EXCP. So; channel program; so the caching was there. I took that code back with me to Poughkeepsie, told them about it, and they said no, the 10 percent thinking again. And I said okay, but this wasn't that much code to do. No, no; extra lines of code and the 10 percent thing, we're not going to do it. Okay.

Yost: Removal of terminal profile?

Worley: Yes, we recognized, back in research, that things other than datasets were important and we felt that terminals were going to be important. In some ways, IBM was right in removing it, because they ended up doing it in a more general way later and probably that was a good thing. There were two different approaches. The RACF database has things called templates, which describe each kind of record. And the terminals were easy to do because you just define a new template and everything was just there. The accessing of this data was done by the RACF manager. Well, it was so general that it was happy with anything that you told it to go get. So basically, it just retrieved terminals. In fact, we never told customers this, but at one point you could actually just define your own templates and it would use them. We never said anything to customers about that, but that code sat out in the field for quite a while. I'm not even sure development realized that; I didn't tell them. Some things you don't tell them.

Yost: Removal of resource consumption information?

Worley: Yes. The original reason; part of the original reason for this code was an inventory to keep track of everything. And at that time CPU time was important, space was important, some people considered lines of print important. So there were parts of the inventory data set, or the RACF data set, that had counters in them or values that you could increment by something or other. And the development people said well, that's not security. And I said it kind of is, but okay, I understand your point. It's not security in this particular sense. So that code all kind of got pulled out and I worked on that later back in

San Jose and basically got it running because at San Jose, we had a huge contention for CPU time there. I mean, you had these high powered people in the molecular chemistry area fighting almost tooth and nail, you know, really get pretty violent. So I convinced my manager, look, I can do this; I can solve this problem. He said okay. But that's why it did not come out; it just wasn't really security. It was, but it wasn't.

Yost: And then you mentioned protection by default. In your opinion, was that the most significant?

Worley: I think it was. I think that was; I think we made a huge mistake there. But to defend data management for a minute, the data management macros are called RAC-something. We wanted to change that. We — development, myself, and Poughkeepsie — wanted to change that. They came back and said that'll cost you three months of test time, which is why there were macros called RAC-D-E-F, because they refused, unless we were willing to agree to a three-month delay, and Poughkeepsie wasn't. So we were stuck with the way data management had done it, which is really the way Gaithersburg had done it. That was just life. They later changed those interfaces so that you could define a profile that didn't have to have the bit on, but it was kind of a silly change. There was a lot of to-ing and fro-ing in order to make the bit useable, because it was just not useable. There were just too many things wrong with it but that early in the game, that was the way that it was done. Looking back, I understand.

Yost: And when you inherited the predecessor of RACF, what was it programmed in and then what was it later, as a product, programmed in?

Worley: Okay. Part of the stuff in what we did in research was actually written in PL/1.

The utilities were all written in PL/1. Hindsight says that was a bad thing to do because PL/1 had a runtime environment and it was priced. If you look at it from a research side, well, who cares but if you look at from a product side, you do care. You don't want to have to drag along something that is priced in order to get this other priced product. So the utilities were all rewritten in PLS, at that point. PLS had many names. BSL, PLS, PLX, PLAS, but they were all written in that and that was a transliterator. It basically took a high level language, converted it to assembler, you ran it through the assembler, and that's where the code came from. Very little was written in assembler, I mean, raw native assembler; very little. It was all written in PLX except the PL/1 stuff, they rewrote that, which turned out to be kind of easy to do because it was in PL/1; one high level language to another high level language and, you know, it was reasonably straightforward to do. So in some ways, that was probably a good decision. If we had written it in assembler; oh, we couldn't have written it in assembler, we wouldn't have. That would've been a waste of time.

Yost: In addition to the survey that you'd heard about, was there any mechanism for customer input into the product?

Worley: Not in this early stage, no. I did present what we had done as a corporate product in a couple of SHARE meetings, and I do not recall which ones. I did present it to some customers at research. I even; it's a Japanese company; can't remember now. But yes, we were encouraged to talk to customers in that way, and we did get some feedback. But that was when it was research stuff. Once it got to development division, I didn't talk to customers. I wasn't allowed to do that; but I did do that because I would go to the SHARE meetings after RACF became public. It kind of was; I was IBM research so what I said did not reflect on development, and I made it very clear that this it strictly my opinions based on some background. And then as customers would come up with their requirements, I would try to clarify them; say, well, have you thought about doing this? And the requirement got refined when it got to development. That was probably not quite ethical, but it worked out really good. And I think most customer knew what I was doing.

Yost: So you could kind of informally help customers along at these SHARE meetings.

Worley: Right. And then when I; after the temporary assignment to Poughkeepsie, I went to England and worked in the Hursley lab. But I did a lot of customer visits, talking about the various releases of RACF. When a new release was going to come out which would've been two or three, release 2 or release 3, I could go talk to them about it; almost to the point of it got announced in the United States, it hadn't been announced in Europe but I could talk about it and that made me popular with the SEs; the IBM SEs. So I talked to a lot of customers, you know, Lloyd's Bank, and that kind of people. And got sometimes some really good feedback. I can't remember any of it, but (pause)

Yost: When the original RACF was released in 1976, do you recall its initial price? How IBM priced the product?

Worley: Ooh, this is a touchy subject. It was expensive. I can't tell you exact numbers because they wouldn't tell me numbers. And then there was a big discussion about ACF2's price. ACF2 came in under RACF and there was a big discussion about what they were going to do. I was never part of those; in fact, I refused to be part of those discussions because I just didn't think it was my position to be involved in that. But the pricing was kind of a sore point with IBM. I think they messed up by making it too expensive. That was stupid. There was some funny things later; I don't know if you know what the SAF interface is.

Yost: No I don't.

Worley: There was a change maybe in; it was a couple years after RACF went out, and it had to do with trying to come up with an interface between the external security manager and the IBM operating system such that you could replace the external security manager, you could see the handwriting on the wall, with something like ACF2. So this SAF interface came out, and there were a couple of people in Poughkeepsie that kind of came up with the idea, and I think it's a pretty good idea. At the time I didn't like it because I knew what they were doing and they would admit to me what they were doing. They were trying to still get our presence in, even though maybe RACF wasn't going to get

used, but the customer's ACF2 was going to get used; but we still had a presence in there. And that's kind of why SAF appeared; at least my understanding. But the price was a bad thing, I think. And it's not unique to IBM, I've seen it elsewhere in other companies where you come out and you try to get as much market share as you can. But if you'd been more reasonable and lowered the purchase price, you probably would've gotten even more. But, didn't happen.

Yost: To what extent did, in your opinion, did not having protection by default hurt the product in the short term?

Worley: I think it cost us an awful lot of money because we had to do an awful lot of things to make the protection; the reason; there was a facility to RACF put in — I don't remember when I put it in — called MODEL that allowed you to allocate this one like that one. And that was put in primarily because of that lousy bit. MODEL was a sort of generic, but not really. So you could say allocate a.b like d.f and it would go pick that model up and use it. But it had lots of flaws and it has the same flaws that ACLs have in UNIX. They have the same flaw. They have these models, and you use them once, but if you change the model, nothing changes. So UNIX security with ACLs has exactly the same problem that MODELS did and somebody didn't look around enough to do; they didn't do their homework and realize whoa, this is not a good thing to do. But MODELS appeared because of the bit; because there was an engineering design group in San Jose that said look, we like the protection but it's just too hard to use. Can we allocate things like other things, automatically. I figured out yes, I can do that. But looking back, I

realize that was not a good thing to do. It was at the time but we were basically trying to get around the unuseability of that bit. We had problems with HSM and the bit, hierarchical storage manager. It just – *bah!* I would bet we probably spent a million bucks fitting that dumb bit in throughout the years. Fortunately, now it's pretty much gone. About the only people that might use it are government people and I can understand why they would.

Yost: When RACF came out do you recall who some of the first customers were? Or what industries?

Worley: The problem I'm having is that I wasn't in the country when it came out. I was in England. In England, who was using it? Banks, Ford Motor Company used it in England, one of the big department store chains I can't remember which one, something like that. That was in England, though. What it did in the United States, I don't know because I was out of the country. I had contact with people but not that kind of contact.

Yost: You talked a bit about your work with customers in the U.K. and Europe. Can you expand on that? What types of things did you help them with in terms of implementation and effective use of RACF?

Worley: Well, in a second assignment, in 1984, I went down to British Telecom a lot because they were using IDMS, and helped them configure what they were doing to use RACF. I spent a couple weeks down there, traveling the train from Cambridge to London

and looking at the code and saying okay, now if you do this this way, that'll work. The first time, it was pretty much; I gave customers presentations on what the facilities were and then took their questions and said okay, this is what it'd do, this is what it won't do. I don't remember any of the customers from Europe. I remember going to Paris one time; I don't remember who the customer was. I can't be really more specific than that. That's a long time ago. (Laughs.)

Yost: To your knowledge, were most of the early customers for RACF commercial customers or was it also important within the federal systems division for government trusted systems?

Worley: From my knowledge, it was commercial but the second time I went to England one of the customers was the U.K. Atomic Energy Group, down near Oxford. GCHQ, or something like that. I can't remember. They were looking at source code. They had access to the RACF source code, they were that kind of customer. They had access to source code; that's the kind of customer we can't even go to the bathroom unescorted. I know. (Laughs.) They were really interesting. I heard rumors that — rumors — that CIA and NSA were using it but don't know. And if they were, they were looking through the code, I'm sure.

Yost: Seeing if there were vulnerabilities?

Worley: They were probably looking for back doors. I got asked that question several times, did you leave back doors in? And the answer is no. Absolutely not. There are no back doors.

Yost: Do you have a sense of how IBM management viewed the competition from ACF, and how that influenced decision making of the RACF product?

Worley: Didn't influence me at all. Two reasons; because I sat out on the West coast and most of that kind of stuff I wouldn't have heard. If I had been a manager — I was a manager at one point — they didn't communicate; interdivision communication in IBM was interesting. Formally, there isn't much. Informally, there is, if you know various people. But I didn't really have much information on ACF2 is doing this so we ought to do it. I wasn't allowed to know that kind of thing; that was forbidden. At least that was what was told to me. So if ACF2 came out with something new, I didn't know about it. Whatever RACF did, it was either the RACF people that thought of it or I thought of it but it had nothing to do with; I never even went to ACF2 presentations at SHARE, if there were any. I didn't go to them because if I had gone to them I would've heard about it.

Yost: When development took over RACF, how big a group was working on it?

Worley: Eight or 10, something like that. And then it kind of deteriorated. I think I might have mentioned that after 1.3 the code kind of stagnated and I had gotten a hold of all of

the customer requirements. I don't know how I got 'em, but I got 'em. So all the customer requirements from SHARE, I got a hold of them. Normally, research doesn't get that kind of thing, but somebody I knew back East gave it to me. And I had some lists of things that the local engineering group wanted, so I kind of came up with a list of things that I could do. I had access of the source from Poughkeepsie; I had a copy of it in San Jose. Don't go there. (Laughs.) I had a copy of it and some people knew I had it. So I started working on this list, then I went to a SHARE meeting; I think it might have been in Toronto. Ron Hankison was the development manager at that time and he's a very; he looks forward to things. I mean, he's farsighted. He's not like most developers that look to the end of their nose and can't see any farther than that. He looked forward and he said, you know, you've got this list of things, what if we put this out as a new point release, like 1.4? I said yes, I can do that. So I was one guy; I got some help from Anne Lescher and Jane Portman in Poughkeepsie, with Hankison's blessing saying come back, we'll pay for your trip, and you can talk to them and we'll come up with a list of things. So I was kind of the designer, the developer, and the tester. One guy.

Yost: Wow.

Worley: So in a year, we did something that had never been done before. Not only designed a release, coded it, and produced it in one year. That had never been done before. Now, it wasn't quite that good. Lot of the testing didn't get as thorough as it should have; there were several bugs that went out in the code and that just; you know, one person; I had no test group. But it rejuvenated RACF to 1.4. Not everything that was

done could get into 1.4. There was a couple things that just; lines of code determine how much goes into 1.4; but 1.4 helped a lot. Password rules came into 1.4. Before that time, I didn't care what kind of password you used, it didn't matter. But the rules came in so you could determine; each installation could determine their own rules for the passwords; how many passwords you kept for history so that you could prevent repeating of passwords; that kind of thing.

Yost: So in release 2 and 3, you were more like a consultant?

Worley: Yes, because I was in England. They'd send me the specs early and I'd read them and make comments back, but by the time you see a spec, it's pretty much; the concrete's beginning to settle.

Yost: Do you recall if there were any suggestions you had?

Worley: I don't recall the suggestions I had. I know I don't remember seeing anything that I felt was bad, although some of it was kind of funny. In 1.3 they added something back that had been taken out; I thought that was amusing, but I can't remember what it was.

Yost: You mentioned some advances with release 4, were there other things that you had to have?

Worley: Passwords were there; the MODEL was there; there was something called Fastpath, which marketing didn't like because that implied a negative performance; we're fixing a performance problem. That got pushed into 1.5 and it became called "Global". I always thought that was funny but I understand, you know, marketing has its rules. So global pushed into 1.5. One-dot-five is when generics appeared; things that Rich Gustki; he and I had talks back and forth on that and there are still some parts that he and I didn't agree on; but from a security aspect, some of what generics did probably wasn't the best. I think it weakened the security, which since then, IBM has had to do some things in order to strengthen it back up but the basic idea was a really good idea. It was kind of like MODELS, but not, so I could have a profile called "A-dot-asterisk" and it covered everything. The asterisk said everything in that qualifier. And that was a good thing; our customers loved that. They really liked that. What else did I do in four? Password. Sounds like a little thing but customers liked it; the last access date appeared. When you went to RACF it returned to a message that says "last time you got on" or "the last time this ID was used" was this date in time. Strangely enough, customers really liked that. I thought oh, that's interesting. That suggestion came from the engineering group in San Jose.

Yost: To detect an unauthorized use.

Worley: Right. The engineering group said, you know, it would be nice if you could do this message. I said, no problem, I can do that; that's easy. Now it caused a problem with CICS but CICS is a very strange environment and they later solved that. But CICS had a

problem; I suspect IMS had the same problem because they kind of log on, but not really; they kind of start a session, but not really. So there was a problem with CICS and that would have been fixed if I'd had more contact with the Hursley group. The Hursley group is an incredibly smart group of CICS people. They are really fun to talk to. They come up with the greatest ideas, or what-if, and oh yeah good idea. I really enjoyed my time at Hursley because I talked to them a lot.

Yost: How was the environment at Hursley different from IBM East coast and IBM West coast?

Worley: Oh, East coast, a huge difference. I mean, East coast development division, development organization, is stiff, I guess that's the way of; lines of code is the all-meaning, all-doing thing that controls it. Hursley; okay, let me back up a minute. When — I don't know if I mentioned it, but — when they wanted to move what I had done in development, they wanted me to move to Poughkeepsie. I said no, research can get away with that. No other division can get away with that but research can get away with that. I said no, I'm not going back there because I've been there. Don't tell me what it's like, I know what it's like. (Laughs.) So my manager said, well, give them something, tell them something. Don't make it such a hard-nosed thing just say I won't go unless you do this. So I said okay, send me to Europe as a temporary assignment. I knew they would never agree to that. Next day, where do you want to go? It was that simple. So I said, England. I had choices of Hursley or some place up near Scotland. Well it turns out Hursley's computer center manager was from Yorktown Research and I knew him. I thought hmm,

there's a no-brainer. I'll go there because I know him. Well it turns out that the guy who wrote HSM, Jim Meyers, was in Hursley. The guy, John Palmer, who I mentioned with the caching, was in Hursley. So it was a no-brainer for me to go there and we were part of the Yanks. We were part of the computer center but we didn't do anything for the computer center, we did our own thing. So that's where I started working on the resource consumption thing, basically that's what I wrote when I was there. You look back on it and you think, no way that could've happened, but it did. It was really incredible. I didn't open my mouth after that about saying something dumb like send me to some place, because I figured there's no way they're going to do that, I'm safe. Little did I know.

Yost: Now the release four, that was the first release of RACF to be evaluated by the National Computer Security Center?

Worley: Yes. At what, C2 level? C1?

Yost: I think you mentioned C1.

Worley: Yes, C1.

Yost: Were you involved in that evaluation process at all?

Worley: No. I only found out about it after the fact. A lot of that kind of stuff I was not involved in. I wish I had been; not from trying to help but trying to learn because part of

research division's function is to learn, and without bias. You never know when something might be useful some day in the future so you should learn. They beat that into me pretty good and I learned it, I think. I'm not sure I answered you about the differences in the environments. England was very laid back. We were the Yanks, the three of us. We usually came in early and stayed late, which gave us a reputation. You know, oh, the Yanks are here. The rest of the people would go home on Fridays about 2:00 in the afternoon; we were usually there until 5:00, mostly because we really liked what we were doing. But everybody was really friendly to us and it was a really nice environment — after I learned to speak English. One of the guys that I got to be very good friends with was cockney. You have no idea how difficult it is to understand somebody who's cockney. He spoke at a hundred miles an hour with all kinds of slang and I'd just sit there and say okay, let me digest this. And he did it on purpose because he knew; but like I said, Poughkeepsie's stiff. San Jose, that's comparing development to research. Research is kind of laid back; research is "it's important to learn", it's important to go to seminars, even if they might not be terribly interesting or what you're interested in, go to them because you might pick up something. So I often went to seminars. But they don't have such things in Poughkeepsie. Hursley, they didn't have such things but I often spent a fair amount of time with the CICS people, just talking to them. Those guys were the ex-compiler writers from PL/1. When PL/1 transferred from Hursley to Palo Alto not everybody went so a lot of the people that stayed in Hursley were compiler writers, really smart people. And, boy, you could tell when you talked to them, they were smart people. To me, I think that was a big advantage that CICS had because they had these really

smart people that were willing to just kind of look at anything. And I think that's one reason why CICS is very successful when you compare it to other things.

Yost: It was huge for IBM long term revenue.

Worley: Yes. Well, if you compare CICS to that other three letter acronym that starts with an "I", CICS was always coming out with new releases and the other one was very slow. But the CICS people just, they were good. They still are.

Yost: And were most or all CICS users using RACF for (pause)

Worley: That kind of evolved because CICS, in the beginning, had its own security rules and that kind of evolved into RACF. Part of that is because of a guy called Phil Emerick, who I often saw in Hursley. He was in — um, where was he? — he was not in development, he was in something like the support centers in the U.S. and he was the one that kind of, I think, pushed CICS people into using RACF. He now works for Vanguard. (Laughs.) He was really good at; he had a very good relationship with CICS people. He also liked the idea of RACF and felt that it was the place to go. I think that his encouragement caused them to eventually get rid of their own security rules and get to RACF. Now you can talk about the other thing, the DB2 thing, in Santa Teresa. They had their own security rules, too, but I kind of think there was pressure implicitly from CICS. Here's CICS and they've gone to RACF, but DB2 has not gone and DB2 has now pretty much gone to RACF. That had to do with some political differences between Santa

Teresa and Poughkeepsie. I got intimately familiar with that. (Laughs.) But DB2 has come around, done their thing, that's good.

Yost: After release four, to what extent were you involved in releases five through nine in the 1984 to 1990 period?

Worley: I read the specs a lot; 1.9 was the B1 security, which affected pretty much everything in the mainframe world; everything in the MVS world because, you know, the base control program got changed; that's where mandatory access control came in and I spent a fair amount of time reviewing the RACF spec and that kind of thing. I don't recall; never got anything involved in the actual base control; what JES is doing or the operating system itself. But it was quite a bit more RACF side because that's when SECLABELs and SECLEVELs came in. I don't know when one-nine was, but it was a long time ago and not very many people used it. But the concepts it put in are beginning to be used now by customers, not only government customers but commercial customers. People are realizing that it's kind of important to have labels on items, and it's kind of important to have levels on security. And it's important to not allow write-down from this level A to level Z, or something, so that you declassify things. Some customers are beginning to realize that this is a good thing. It's not easy to get to that level of security, it's difficult. Lots of things you took for granted to do you no longer can do; you just can't do them anymore. But once you get used to it, it's not bad.

Yost: Was there a sense of, when the National Security Center came out with the Orange Book first in 1983 and revised in 1985, how this was going to impact how you would interact with customers and what specs would be written?

Worley: It didn't impact me. I interacted pretty much the same as I always did because I could claim I was in research and what I say is not a commitment on anybody's part. So I pretty much could; well, I couldn't tell them some things, though. I can't tell them future plans; I can't tell them that; I know that. But I can tell them, if something just got announced, I can say yes, it's this feature and it does it this way. That's perfectly legitimate.

Yost: In 1988 you were one of the authors of a paper entitled, "An Access Control Model: If We Only Had a Clean Slate." Can you discuss the context of this and what lessons were learned in looking back?

Worley: Access control lists are very useful but they're very cumbersome at times. There's always been a decision in RACF that I thought wasn't good. We allowed group names in access control lists and the idea was that if the group's in it, therefore all members of the group; if you're a member of this group, and the group is in the access list, and we can tell that you're a member of the group then you can do whatever the group says. And I always felt [that was a] bad decision because that implies that you should have protection on the group membership, and you really didn't; at least not to the point that it impacted or was known to the access list of the particular resource that the

ACL was part of. Even today, there are people that say just put a group in the access list, it solves the whole problem. It doesn't, it creates a whole new problem. And it's hard for me to convince people that no, no, no, no, this is a bad thing. But I guess I look back and I say, well, okay. Maybe I'm just too early in this whole process. I think it was things like that. ACLs in RACF terms got perverted where you could do this if you were going at this data set on that volume. And the ACLs became very complicated, and it was because we were trying to stay within the conventions of the current ACL structure and it was wrong. It just seemed like it would be easier to say we screwed up, let's go back and do this right. But that isn't the kind of thing development's ever going to do. It's the kind of thing research would do, and I think that's what the paper was getting at. I don't think it came out and said exactly that we screwed up, but it said there are these other things that we could do that maybe would be simpler, rather than this complicated world because the way access lists are now, they are so complicated. It is so easy to break into people's systems. All you gotta do is look at their access lists. If you have access to their access lists, you can figure out a way in that's legitimate. They don't want you to be in, but you can get in and most people don't realize that. They say how can you get to an access list? No problem, they're in memory. Well, where are they? No problem, they're in a data space. You got the data space protected? I didn't know you could do that. Yes you can protect it. If you don't protect it then I can dump that data space, and all the profiles, and all the access lists are right there. And you don't even have to be a rocket scientist to read it, they're just there. And it's that kind of thing that we did wrong. We just didn't; I don't think we understood the implications of computer security. We didn't understand how big it really is.

Yost: You were an active participant with SHARE for a number of years. Did you also attend any conferences in computer security?

Worley: Usually not, no. I wish I had, but no, I didn't. I was in a funny position because I was a member of a security architecture board inside IBM. That was not known outside the world, outside IBM, that it existed. We often saw things that were what you'd call highly IBM confidential. So you get into this situation that I know this, but how much can I say? I remember there was a product back East that had this huge hole in it, and in there a while. I won't tell you what product. We, the board, strongly recommended it get fixed and the reason it hadn't been fixed was because of lines of code. Lines of code?! This is security! You gotta fix it! And, no.

Yost: Without talking about anything that you shouldn't, could you give me some idea of what the mandate or the goals for the security architecture board were?

Worley: It was pretty much to; we saw things that were either in plan, or thought about, and then we kind of came up with a set of recommendations that said you should look at this, you should look at this, perhaps you ought to avoid doing this. We tried to put everything in a non-confrontational way. The one thing you do with developers is you don't want to confront them upfront and say that's a dumb thing to do. You're going to lose right there. And so we tried to put everything in kind of general ways, saying well, if

you did this then this is going to happen, or could happen, and that kind of thing. How much ever was acted on? Don't know. We never really got any kind of feedback on that.

Yost: Do you know when this board was formed and were you one of the original members, or did you come on later?

Worley: I was kind of one of the original members. There was maybe 15 members, kind of from all divisions. There was member from RACF development; that was Rich Guski. And I was there; I'm not sure why I was there. I was there because of history. The main organizer of the board knew me and called me up one day and said would you like to be a member of this? And I said sure. And he said, oh by the way, we'll pay for your trips. Well, that's always a good deal with research. So I was kind of a member. When did it start? Sometime after 1976; maybe up through; I don't know when it ended because I went to England, again, in 1984, so I was not involved in it after that. So I don't know how long it went; I don't think it went very long after that because I kind of think it; I don't think it was as effective as they wanted it to be. Pretty hard to be effective when you're trying to deal with product schedules. They just don't want to listen.

Yost: In the time when you were in the U.K., in the document that you provided me, you indicated that Cambridge University was a major test run implementation, that you helped to do.

Worley: Yes.

Yost: . . . the RSC system. Can you provide some context and elaborate on that?

Worley: They came out to San Jose in 1983, something like that, and I presented resource control to them. Now this is the thing I had started working on in 1976, kind of off and on, and it was the resource consumption thing. By 1983, it was running; as much in production as anything on a research machine. And they had done something similar, not quite as elaborate, and they were interested in what I had done so I gave them a presentation on this. And there were two IBM guys — can't remember their names — and we went to lunch over at the country club in the San Jose site, with the Cambridge guys and I said something about well, I had been in England in 1976. I said I really liked it. One of the IBM guys, the SE, sort of jokingly said would you be interested in coming back? I said sure. And he looked at me and he said well, I'll see what we can do. I thought, all right, here we go again; because I figured I'd never hear anything more about it. But not too long later, maybe a week or so, my manager told me, hey, Cambridge is interested in you coming there for an assignment. My manager was not happy about that for whatever reason, I don't know. I never asked him. But I said okay, yes, I'd be happy to go because then I could do this RSC in a real customer world and see how to fit it into what they had done. That presented some interesting challenges. It was written in PLX, and they didn't have a PLX compiler. And they never did. I overcame it, somehow, I figured out a way of getting around it so I could put the code on their system. And then I took their idea; they had something called "aging." They could measure resource consumption, and then if you didn't use a resource, your consumption aged down, which

was a really neat concept. Basically, I was able to do that by just sticking an exit in my code, and do the aging algorithm, and then put the number out there. And it worked out pretty well. Then we came to the problem of okay, now how can they get this? It got into IBM politics. So I kind of did that for maybe a year at Cambridge, and I spent my second year of my assignment I spent down in the branch office talking to customers, because we just couldn't go any further. Cambridge wanted support, while research doesn't support things. And they wanted it to somehow get into the product line. Well, that got into a whole political game between Poughkeepsie and Santa Teresa, and I tried that after I got back from the U.K. and there were "some" politics I just couldn't believe. There were some people that just didn't like each other, therefore, whatever they said, the other guy said they weren't going to do. Kind of reminded you of Boehner and Obama, but it was that kind of thing, you know? I'm not going to do anything you want. So I got really close to getting RSC out as a product but the two divisions just couldn't agree. So in the meantime, I had talked to Poughkeepsie and they realized that gee, this RSC is an address space and they'd been looking for a place to put some of the RACF stuff. This was Rich Guski, and we talked about it for a while, and he said well, how hard would it be to pull some of this stuff out and to just kind of leave a skeleton address space? I said, that would be easy. So somehow, it got into the product schedule as 1.9.2, I think; a RACF address. But it really didn't do anything. It was a presence more than anything. That was the comment we got from customers; what good is this thing? It did two commands but nobody ever used those commands, so it didn't do really anything. But it offered a home, and that's what we were looking for; it offered a presence. And it's very hard to tell customers this was just; watch this space. I could say that as research; development

couldn't say that. And I said it as research so we got the address space out there. The address space then became a home for RACF remote sharing, which Rich Guski and I talked about, and I ended up doing it out in San Jose. It transferred to Poughkeepsie, and then they put it out, I don't know, a year or two later. So I got fingers in 1.9.2, little bit in 2.1, and I don't know when remote sharing came out; 2.2 or something, I don't know. So I had fingers in there and now we have a RACF address space. They haven't gone near as far as I believe they should but there was a person who was very anti-address space in development and he felt what if it goes down? And I said, well, you write code that doesn't go down. It's not impossible. He said, we have no address space like that. Well it turns out we did; it's called catalog address space and all the catalog work is done in catalog address space. If it goes down, you are in deep trouble. So why can't we do this with RACF? It's very hard to argue with somebody if he's on the East coast and you're on the West coast. You're not going to win. So they haven't done anything with it so far.

Yost: That sounds difficult. But Guski was in development on the East coast and that collaboration worked.

Worley: Yes it did. It worked really well, yes. Rich and I tended to think a little farther out. The other developers, like I said, have this end-of-nose attitude, and I understand that. But you can't really do that otherwise products stagnate.

Yost: So it was in 1992 that you officially retired from IBM but you continued to consult?

Worley: Yes. Actually, I tried to retire a year earlier and the senior level manager above me wouldn't let me because I was bringing in a lot of money per year, because of my development activities. Development was paying me to do development work and he didn't want to give that money up. He found out all kinds of reasons why I couldn't retire. So I waited a year and tried it again. Fortunately, the head of personnel, Human Relations, now, was a good friend of mine and she helped me. (Laughs.) So I did get retired then but prior to that time, I had become friends with the owner of Vanguard. He had been in an IBM presentation at Poughkeepsie where I had presented something. He said look, I'll give you a job, whatever you want to do. I said well, I'd have to retire first. Well, retire. So that's kind of how it happened. And we wanted to come back up here so my criteria with him was look, I'm going to retire, I'm going to move to Oregon, I have to be able to work remotely. No problem, so here I am.

Yost: Has this been continuing to consult on the same types of things or completely different?

Worley: No, once you're out of IBM, you're out of IBM. It's sort of like — I was going to say the CIA, but that's not true — once CIA, always CIA. No, in fact I'm very careful what I say. You've noticed there are times I won't answer things. I like my retirement pay.

Yost: Right. So interaction with RACF is over?

Worley: Only when I happen to see them at; the company I work for gives a security conference every year and I see Mark Nelson there, who I know. Rick Guski retired from IBM and he works for Vanguard as a consultant-type thing, so I see him, and talk; but I can't communicate direct to IBM people that same way. I'm an outsider now, I can't; I just can't do that. I don't think it's ethical and that's, I guess, my hang-up.

Yost: And finally, are there any questions I didn't ask, topics I didn't cover that you think are important to understanding the history of RACF and its place in the history of computer security?

Worley: I remember giving a presentation at one of the Vanguard conferences, and it was one of those big general things. I remember looking out at the audience of about 300 people, RACF users, and in the back of mind I thought, what have I done? You never really understand when you're doing something, the implications of it. I guess that's good; it's probably a good thing. There's some really good people I worked with; really good people. Some of it was just purely because they were so good, it was their attitudes that give you incentive. The same research division or department I was in, John Backus was in. FORTRAN leader. Ted Codd was there; relational data bases. Those are some really high-powered people. The guys that developed SQL were in that same department and I always felt very, very honored to even have been in the same building with those people. But when you go to seminars with those people you hear attitudes that are really

great. They are interested in learning, they are interested in things. It's not focused down on a little minute thing, it's this big interest. I don't know how I convey it to you.

Yost: Did IBM Research change significantly in your opinion during the time you were there?

Worley: It did. By the time I was there, at the end, it got more product oriented, and I think that's okay.

Yost: As I remember, Ted Codd went on record saying that's why he left.

Worley: It became difficult. It wasn't as bad as development. Development was much worse, in my opinion. The group's manager might have 20, 30 people. You can't be a manager with 20 or 30 people, you just can't. I was a manager several times in research, and at most, one time, I was the second line manager temporarily and I had something like 20; that nearly drove me nuts. Most of the time I had five or six. That was okay, I could deal with that. But it IBM Research changed a lot and it got difficult. I wouldn't want to go back. Just don't think the same attitude is there. And that's okay; times change and that's the way it is. I've often thought it might be nice to go to a university but I'm not sure I'd want to do that either. I like sitting here all by myself. I had an office in San Jose; normally, the door was shut and the blinds were drawn. And there was one little light on. People knew they could knock and I'd be more than happy to have them come in, but I like the solitude and that's, you know, I could stare at a board and draw syntax

diagrams. And it was good; I liked that; I miss that now. I kind of just do development work now, and some of the work is right here. Managers leave me alone. They pretty much leave me alone, but still. If I had to go back what I would do is probably work on the RACF address space and really make it an address space because I think you could do such a thing and I think it would be a big advantage. We screwed up in the past, not only on the RACF, but IBM making things external, such that, once it's an external a customer depends on it and it's really hard to evolve things then. When I started doing some of the stuff for remote sharing, I found stuff in there that was RACF 1.3 level that had no meaning anymore and I pulled it. Never told them that, but I pulled it. That happens, and I understand that happens in development, but at some point it's kind of like that paper, you need to start over. (Laughs.)

Yost: Well thanks so much. This has been fascinating and a terrific help to the project.

Worley: I hope so.