An Interview with

DAVID ELLIOTT BELL

OH 411

Conducted by Jeffrey R. Yost

on

24 September 2012

Computer Security History Project

Reston, Virginia

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

David Elliott Bell Interview

24 September 2012

Oral History 411

Abstract

David Elliott Bell is a mathematician and computer security pioneer who co-developed the highly influential Bell-LaPadula security model. This interview discusses the context of his pivotal computer security work at MITRE Corporation, and his later contributions at the National Security Agency and Trusted Information Systems (including his leadership on TIS's Trusted Xenix B2-rated system).

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 Charles Babbage Institute, and I'm here this morning on September 24, 2012, in Reston, Virginia, with David E. Bell. This is an interview for CBI's NSF-sponsored project, "Building an Infrastructure for Computer Security History." David, I'd like to begin with just some brief biographical questions. Can you tell me when and where you were born, and where you grew up?

Bell: I was born in 1945 in Liberal, Kansas. Grew up in North Carolina — moved there before I was two — Winston-Salem, North Carolina.

Yost: And in elementary, middle school, and/or high school, did you have particular interest or strong aptitude for mathematics?

Bell: Yes. My school system was in transition so I went from elementary school straight to high school. I was sent to high school for five years. Yes, math; always was good with math.

Yost: At that time in your life, during school, so pre-college, who were the greatest influences in your life?

Bell: I don't know that there's anyone beyond my family. A minister at my church, I suspect, and possibly the scout leaders. So probably that.

Yost: What year did you begin your studies at Davidson?

Bell: Started at Davidson right after I graduated in 1963, just before Kennedy was assassinated. I was there four years, to 1967.

Yost: Did you start off as a math major?

Bell: The first couple of years I kept changing my mind about what I thought I wanted to do. The first year I wasn't exactly sure; then I thought I would do psychology. That seemed like a bad idea. Then I decided second year I would do physics, but I was a year behind in getting all of the labs and so on done and I went to mathematics as the easy choice. So I came to mathematics because that was the easy choice after everything else that I thought about didn't pan out.

Yost: Okay. And when you were in college were you thinking about any particular career aspirations or career path with your mathematics major?

Bell: No.

[Interrupted by phone ringing. Break in tape.]

Bell: You should probably repeat whatever you just said because I've forgotten.

Yost: You were talking about your choice of majors fluctuating early on, and then I asked if you had an idea of what you wanted to do for a career with your math major while you were in college.

Bell: While I was at Davidson — you know, the prototypical liberal arts school — I concentrated on getting educated. And I wasn’t really thinking about where I was going and what I was going to do afterwards. My father was trained as a music teacher at a high school. He was the choir director. My mom played the piano and then the organ. I was in the choir. I was in the band. I was in the orchestra. So while I was there at Davidson, I took the full liberal arts requirements including Bible, English, and history, and I took an advanced class in music theory and history . . . analysis and, yes, history. Intending to get myself educated so that no matter what I did, I would be in a position to learn the new stuff because what was current was not going to stay current and one needed to learn all of these things. So I really had no notion of what I was going to do for a career afterwards. At my senior year, I applied to graduate school, but I also went searching for jobs and I got a job offer at Edgewood Arsenal in Maryland to do God-knows-what. I don’t know what they intended for me to do in Arsenal. So I went off to graduate school because they were going to pay for it and it was an easier thing to do than going off . . . . And it was 1967, and there was the risk of being drafted, even though I was married by that time.

Yost: Did you have any exposure to computers at Davidson?

Bell: I did, but not a great deal. I took one class in numerical methods in the math department and we had a small IBM computer, whose number I cannot recall any longer, that we could make use of for doing things. Some of my friends did their physics labs by writing a quick program and running over and printing it out. I wasn’t so good, but I did my numerical methods there — this was punched cards — and you could write programs that allowed you to say, “keep doing this loop until switch 9.” And when you were finished having it run, or your time was out, you would flip switch 9 and it would terminate. So that was the extent we used FORTRAN. That was all we had for computers; there were no computer classes. In fact, the first year in graduate school the cover story on the *Journal of the American Mathematical Society* was “What to Do Till the Computer Scientist Arrives,” because there was no such thing in 1967, 1968. At least in the broad sense.

Yost: Yes, I think Purdue and only a few other places had CS departments by that time.

Bell: Well, what the article talked about was the fact that engineering schools…. In engineering schools, the computer department grew out of engineering. And in business schools, where big business, . . . then computers came out of business. And there are some places where they came out of mathematics departments. Particularly ones that were strong into numerical methods, and so on.

Yost: Right. And so you went to graduate school at Vanderbilt?

Bell: Yes.

Yost: And what were your…. Can you go through your evolving research interests while you were in grad school?

Bell: I was there four years, and when I first arrived there was a standard set of classes that you needed to take, no matter what you were trying to do, which exposed you to algebra, and analysis, and also topology, and a variety of things like that. So that was more like a survey. Unlike some of my classmates, I decided I was going to write a master's thesis and I had the good fortune to be given the Distinguished Professor of Mathematics as my advisor, Bjarni Jónsson, who is an algebraist. And what he suggested for me was looking at consequences of the Feit-Thompson Theorem, which had to do with proving a long-standing conjecture that every — let me see, how'd it go — that every simple group has odd order. The proof consisted of one entire issue of *The Pacific Journal of Mathematics*. And what he taught me, in addition to the material — well, most of the material I found and mastered myself — but what he taught me was how to write a mathematical argument in a way that was easy for other people to follow, as well as yourself. His method was usually to say, "The general result follows if I can prove the following three things; let me show you how. One, two, three leads to conclusion." And then he would proceed. "Let me show you A; let me show you B; let me show you C." So that happened in the first year and a half of graduate school. After that, when I passed the prelims…. There were six of us that passed the prelims at the same time. Five of us went off to Charles Megibben to ask him to be our advisor — five out of six. He also was an algebraist from the Alabama school and we each of us had an appointment with him to see what we were going to do. And when I got to him, he said, "you've been here the shortest period of time; you’re last. Go read.” And he was in the process of moving from algebra to category theory and suggested that I look at groups and start looking at category theory. So I did that and at one point I thought I had solved a long-unsolved problem. Let’s see — what was it called? The Koethe Conjecture. I’d even have to go look up how to spell that one. And one Friday, I gave him a paper after the class that we’d had, saying I’d like for him to look at the paper I had written. He said “Okay, what’s it about?” And I told him and he said, “That would solve the Koethe Conjecture.” And I said “Yes, that’s why I wanted you to look at it.” He called me Saturday to say I need to start preparing a paper for publication and find all the relevant literature. By Monday morning, everybody had heard. Everybody was excited. And Tuesday, he called in son [Ed. Peter M.] Neumann, who was visiting professor that year, and he was outlining my proof. And Professor Neumann said, “Why does that have to be a transfinite number? Why doesn’t it work with three?” My multiplication didn’t distribute; as a result, everything I had done was absolutely wrong. Dr. Megibben never got anything again where I left out one step. I ended up generalizing results from torsion and torsion-free groups in category theory. So I went off to…. When I went to my defense, most of the people on my committee were not algebraists. Dr. Jónsson was, and that was a concern. But he didn’t ask me any really difficult questions. So, I think that’s it.

Yost: And can you talk briefly about your experience with using computers and any programming that you did while in grad school?

Bell: In graduate school there was no programming. The summer before I started graduate school, I had gotten a job with the National [Ed. really the “American”] Association for State and Local History, who had a grant. They had gathered a large amount of information, which they wanted to essentially sort and they needed somebody to do it. And I ended up being . . . . And doing all of that with sorting machines, as opposed to having to write programs so computers didn’t really come up. The biggest thing, when it came time for my dissertation, I couldn’t afford to have a typist do it. So I made arrangements to use the IBM Selectrics and when the secretary left at 5:00, I arrived and I typed. I needed three different type balls. In category theory you draw diagrams — a symbol here, then a straight arrow pointing right, then another symbol there, and then arrows pointing down, and sometimes diagonal symbols. And so I had to type, type, type, type, type and later draw in the arrows. And I would work from 5:00 until 8:00 the following morning. And I did that for weeks. Had to buy my own type font in order to have one for…. We had one for symbols but we didn’t have one for the particular characters I needed. But no computers.

Yost: It was in 1971 that you finished and defended your dissertation?

Bell: Yes.

Yost: And that was the same year that you started at MITRE? Was that your first job out of graduate school?

Bell: Yes.

Yost: How did you learn about the position? How did that come about that you came to MITRE?

Bell: I went to Vanderbilt, the placement office, and came to them and said I'm looking for a job. They had nothing to tell me because I was a graduate student, not an undergraduate, so they handed me a book and said "you can look at this because these are companies and they tell you what you care about." This was in the middle of what we thought then was a bad economic time. Companies had fallen back. We were in what people were calling stagflation at the time, and all sorts of people just were not getting jobs. I found out later that MITRE had had its first RIF [Ed. Reduction in Force] a couple of years before I did this. But what I did was I took my book — about the size of a medium-size phone book — and I went through and found the people who wanted mathematics Ph.D.s. And I had a typewriter, a manual typewriter, at home, and I alternated weeks. One week I would write. I would type five letters and mail them every day. The next week I spent that time studying. And I did this for months. And one day while I was teaching — the last three years I taught for extra money — the secretary, the one that I had used her Selectric, came to where I was teaching and said, "you've got a phone call." It was a different building. I had to go over to the building. And Brown — what was his name? I'll remember it later [Ed. it was Fred Brown] — had called me up, saying that a telegram was on the way to me that they wanted to have me come interview. So I had sent them a letter, cold call, and they said that they wanted me to come for an interview. My wife, my ex-wife, said, “What is this company?” I said, “Well I don’t know.” I had the information, which was Bedford, Massachusetts. She said, “Where’s that?” So I got out the atlas and I said, “It’s an inch from Boston.” I agreed to go out for the interview, and cut my hair, and put my best suit on and went up there for the interview.

Yost: So at that time you didn’t know it was outgrowth of SAGE or any of its history?

Bell: No. I mean, I had this sort of squib you’d put into the thing: “it’s a not-for-profit, system engineering corporation, mumble, mumble.” I probably hadn’t read that part, all I saw was “mathematics, Ph.D.”

Yost: And when you arrived, what was your first impression at MITRE when you interviewed?

Bell: Well, it was . . . . The buildings are their own, but it was what you would now call an industrial park. It’s more of a campus kind of thing. The parking is off to the side and then there’s walkways in between. My first impression was it wasn’t like a college, I guess that was the main thing. And I was trying to…. I remember one of the things that I said after I went home was that as I walked from place to place where they were having me interviewed, I saw people with sandals on. I saw people with really long hair. I saw people that wore button-down shirts. And I thought to myself, if they can accommodate all of those people then maybe I’ll be okay.

Yost: When they interviewed you, were they thinking of you for a contract that was rising out of the work that the Anderson committee was doing, and that Steve Lipner was on? Or was it just a more general mathematician position; they'd find contract work for you?

Bell: I don't know for sure. I'd have to deduce it. The way MITRE interviewed at the time — it's gotten more structured — the way they interviewed at the time was that the resumes were passed around through the organization and if anybody within that organization was interested, they would sort of sign up. And then within an organization, like a department, they'd look at it and they'd say who should we have look at it? Interview this person. So what appears to have been the case is that they were looking for people that they could make use of. When I first arrived, I was there a year before I started working on this kind of security. I had a half-time job doing a different kind of security, something that actually was a kind of program correctness. They had a prototype and the question was, is this doing the right thing? So I read a lot of proof of correctness for a year and tried to figure out what I could do about their design. So it was not specifically for Steve Lipner's project.

Yost: Okay. And so that was half of what you were doing. What was the other half, that first year?

Bell: One of the people who interviewed me was Len LaPadula, and in fact, he got interested enough that he took me off to show me a lab and we got off schedule and they didn't know where I was. I ended up working with him and being his office mate on something that we called "harmonious cooperation," which was a kind of sharing, an extension of readers/writers. Readers/writers mostly concerned itself with blocking kind of operations at a gross level. What we were doing was looking at narrower sharing mechanisms at a very theoretical level, like the *Communications of the ACM* would be interested in, in such a way to would avoid blocking but also to say "No deadly embrace! No indefinite delay!" — things like that. So that's what we worked on for a year. We thought that was going to make our name and we would be famous. But no.

Yost: And was Leonard LaPadula's background also as a Ph.D. mathematician?

Bell: No. I'm not sure what his undergraduate degree was in but he had worked in software as programmer/software engineer. He had worked at MITRE before. He'd gone to two other companies; had come back to MITRE. Honeywell was one of them; the other was a smaller company and I don't remember the name of it. [Ed. Teledyne] But he was mostly a software type or an engineer, and later he got his professional engineering degree, which is a Ph.D. equivalent — at Northeastern, I think.

Yost: And how did it come about that you… that the two of you began work on, you both, begin work on mathematical models for computer security?

Bell: Well, in thinking about this interview today, one of the things I thought of is the fact that in writing the paper, my 2005 paper, what I did was synthesize much of what was going on, some of which I didn't know. And my experience has been that when, in the computer security field, if you ask somebody about what happened, a number of participants, you get different stories, which may or may not be complementary. Sometimes they're contradictory. So what I have understood looking back, is because of the Anderson Report, a wide variety of endeavors or initiatives were started by the Air Force, which was then executive agent for computer security in the DOD, meaning they were supposed to take the lead and figure this out for the whole DOD. One of the things that they didn't quite understand, Roger just recently told me, that when the Anderson Report said what we've done in the past doesn't work, penetrate and patch doesn't work, Tiger Teams show penetrate and patch doesn't work. They came up with an idea that what you needed was . . . a conceptual design of a general computer system that would be secure. And there were a number of people who didn't think it was possible. They didn't know what it was. Well, every year, the Air Force came to MITRE and said these are the things we want you to do, and they would lay out projects, and within projects, tasks. So out of that came a description and Len and I got called into somebody's office, probably Lipner's, who was in a different department. So we went over there for him to kind of make a pitch. We were actually…. In those days, you got some choice in which . . . what you got to work on. Somebody would say I'd like you to work on this, and you would agree or not. Particularly outside your department. Your boss could say "this is what you're working on." If he was being nice to you he'd say "which of these two would you like to work on?" So somebody made a pitch, probably Lipner, and said this is what we want you to do. We had just done reader/writer sharing. We had proved what we thought was a general theorem about what you needed to know in terms of prior knowledge to avoid deadlock. And we…. So he showed us the description of the task and the task was: take a year, write a report, this is the name of the report, and produce a mathematical model of computer security. Well, we didn't know — *they* didn't know — what that meant. And *we* didn't know what that meant. And we walked back to our office, we said "That sounds pretty boring. Not nearly as exciting as what we've just been doing." So that's how we got into it, but there wasn't much of anything that was more enticing so that's what we ended up doing. I was talking to my wife, who's in the same field, and I reminded her that that first year, in twelve months they dedicated one-and-a-half staff years to writing a model, and in the process of that…. Well, let me see, that's one-and-a-half staff years that year. The next year; I can't remember when Len went off to NASA, the second year. So it was either three or four staff years, over two years. We wrote three reports, not two, and about six white papers. She observed, [that] nobody gets that kind of money for that kind of time for doing those kinds of things anymore. But we agreed to go off and do it and started trying to figure out what we could build on, and there just wasn't much. So when we found that the things in the literature that we could find weren't helping us, and for the reasons that I have mentioned before, some of them were extremely specific about this computer system, ADEPT-50. The paper was about the ADEPT-50 operating system and it was kind of abstracted away from the operating system. Some of the others were at the level of the theory of computation and were extremely general; you knew you could put a computer under it, but it was a long jump. And we viewed it that what we needed to do was to be able to address any computer system so we had to hit that balance that would be not too specific and not too general because we figured you needed a tool. We'd seen what we thought they needed was a way of analyzing, addressing, assessing a real computer system. So we went off to try to figure that out.

Yost: Did you have any involvement in that first year with the people from the Anderson Committee other than Steve Lipner, who was at MITRE?

Bell: No. When the report was finally published, I received a copy but I didn't read it. I mean, I kind of looked at it and I said that's not helping me. So I never actually read it very well until the early 1980s when I went to NSA. So the people that…. So I'd . . . .

Yost: Did you read Ware's Defense Science Board Report?

Bell: No. I became aware of that sometime in the mid-1970s when somebody made reference to the picture, I went off to try to find the picture [Ed. Figure 3, "Computer Network Vulnerabilities"]. And I had a copy for a while. I had a large numbers of copies of things that I had sort of looked at the abstract and then filed away. But I hadn't read it. Somebody had a grand scheme of things that they thought needed to be done, hoping for success in enough places that things would work. And we were one of them. But I didn't know the grand plan. This is the blind men with the elephant. You know, I knew the part I knew; other people knew the part they knew. And our notion of what's going on differs, just based on what we saw.

Yost: On this project you were reporting to Steve Lipner, is that correct?

Bell: Yes. And he was in a different department. At MITRE we called that “soft shell.” You were outside the protection of your own organization so somebody was directing your work but it wasn’t your boss. It was your project leader but not your boss. So, yes, Lipner was running the project and there were a number of tasks that people were doing.

Yost: Was this 100 percent of your time or a lesser percentage?

Bell: The first year, it was…. You know, I don’t remember. The first year I was half and half. The next year . . . . I think I misspoke before; I think we had two staff years that first year. Both of us worked on it full time, I think. By the next year, sometime in the next year, Len went off to NASA down in Houston, and then it was just me. So.

Yost: Can you talk about some of the things you explored to try and get a handle on this problem? What in that first year did you consider?

Bell: You mean the background reading?

Yost: Yes, the background reading.

Bell: Well, we read the ADEPT-50, and we read the Denning paper. Some of the Multics papers we read, but later. I don’t believe that we had collected the full set of the *Fall Joint 1969 Computer Conference* papers when we were doing the first things. The main thing that we observed was that we were reading about where people were talking about protection or security, what they were talking about seemed to vary all over the map. So the first thing we thought we really needed to do was make a very clear definition of what we were investigating because we had also seen people criticizing each other, saying “You say you’re secure but you didn’t do ‘X’”. What we wanted to do was to say, “As far as we’re concerned, ‘security’ is X, Y, and Z.” And if somebody says “You didn’t do W,” we say “Yes, we said that. We’re only doing X, Y, and Z and if you want to do W, then fine, go ahead and do some W but we’re X, Y, Z.” So the main thing we wanted to do was to, early on, was to figure out what our definition was. And in order to have a definition it was necessary to figure out what we needed and what kind of framework would make it “mathematical” or “conceptual”. So we started thinking about that and the notion of subjects and objects was around everywhere so we used the same words but in many cases with less connotation. Ours was kind of abstract. So we went about trying to figure out how to express things. We needed to be able to say there are active things that would like to get to passive things — subjects and objects. And we need something to say we’re going to make a choice, so we have something in the middle. Then we needed to be able to say things change over time, so we had to have a sequence. And you needed to know what your current state of the world is. So you’ve got a representation of what the rules and limits are, but also the current situation.

So we started setting something up like that, that would give us the ability to make definitions. And then we were asked to come to a project meeting and I'm not positive who was there from the Air Force, but Roger Schell was. I believe he was a major at the time. So they went around for each of the tasks, seeing what you were doing. And one of the things that we told them that we were doing…. I mean, the previous year we had worked on harmonious cooperation, and sharing, and deadlock, and so it came our turn and we said "We're building this definitional structure and there's some really interesting stuff. If you reclassify something, you know there are good reasons [not] to delay the classification, but unfortunately that allows someone to deduce something. If you delay, it's possible to get into an indefinite delay." And we were all excited because it was very complex, and it was also something we were used to. Roger Schell gave us direction to assume that nothing ever changed. So we were a little dumbfounded and on the way back to the office we were a little steamed and we said, "There's nothing interesting left; it's boring. There's nothing to be done." So we took what we had done and put it into volume I. Just to make clear how annoyed we were, we called our result the "Basic Security Theorem" as compared to the "Complex and Extremely Sophisticated Security Theorem" that we thought that we were going to be able do. Now, at the end of the paper, we snuck some of that stuff back in, but on the surface we said "Yes, we're being compliant."

Yost: Did Schell explain his reasoning?

Bell: No, he was the Air Force officer. He was also an MIT Ph.D. But he didn't explain. He was right but I didn't think so at the time. You know, 20 years later we were still struggling with this, and in real life. I mean, I had a SECRET Clearance at the time, it was *pro forma*. You don't automatically change things, security classifications. There are procedures you go through, and downgrading in particular is hard. Upgrading, you don't tell people. You don't change the classification of a *New York Times* article, you take a copy of it and you put it in a SECRET folder is what you do. [Ed. Like the Pentagon Papers.] So this notion of changing live was not a front burner issue at the time and it didn't become one for some time. It's kind of a specialized issue of how do you sanitize a release? So we got that out of our systems. Now, there was a minor glitch that came from that. We decided that we had been too forthcoming. We let people know what we were doing so they were giving us direction we didn't want. So we decided we would learn from our experience and we would stay quiet until we had something finished. It turned out that worked to my detriment. I arrived in September the first year, and reviews and salary increases are in January. Somewhat to my surprise, I got a raise three months after I got there. The following year, in the fall, we were being very quiet and as they were preparing the salary reviews they looked at it and said these people aren't producing much of anything, because we were being quiet. So they put in the paperwork and I didn't get a good review. And about early December, my department head and Steve Lipner's department head — and Steve Lipner — told us "You need to make a presentation to these people about what you've accomplished." So we went in and we made our presentation and everyone suddenly got agitated. My boss had already put in for the kind of review that is a suggestion that you better shape up or we're going to fire you, and he told me that he had tried to recall that but that he wasn't allowed to. Ten years later he told me that ever since that time his intent was to pay me so much money I could never leave. (Laughs.) But the first year, because we were trying to keep quiet, it looked like we weren't doing anything. So after we got that first one out of our system; I should mention (pause)…

Yost: Your motivation to keep quiet was that because you wanted things solidified and perfected before you would present?

Bell: The word wasn't in use, but we didn't want to be micromanaged. We went off and said these are the interesting areas we're looking at, and they said (we didn't think with reason) "ignore those." Well, if we found something interesting we wanted to look at it. So we decided that if they didn't know about it they couldn't tell us not to. So ask forgiveness not permission. But after we did the first one we immediately launched into the second one. I should point out that Len had six or seven years' experience over me, something like that, so he was definitely the lead on our task. It came time to publish the article, he had every right to put his name first and what he said was, "We're going to be alphabetical and we'll alternate." So if he had taken his perk, it would've been the "LaPadula-Bell model." I think "Bell-LaPadula" flows nicer but it could've been the other. So for the second one, what we decided to do was to start making it look more like a computer system and add things like access modes and that sort of thing.

Yost: You brought up ADEPT-50, and was the high water mark policy access, at all, an early anticipation of the *-property [Ed. read "star-property"] for you or Len?

Bell: No. I mean, we read it and we disregarded most of it because it was so specific. We did not extract the principle and other people who were critiquing it pointed out that what ended up happening was that security levels started migrating and eventually, at the top, everything was classified System High and that wasn't the original intent. But in the first volume that we wrote, in a fit of pique, we viewed data repositories like a book at a lending library. It was just a whole book. And in the second volume, when we started adding things like access modes — you know viewing something is different than altering it, and so on. That's when, suddenly, we got the idea that we couldn't think of it as books in a lending library anymore, that if you checked out a book and you saw on the edge as . . . . Sometimes people that know things first hand may read something in the book. At Davidson, there were some books that are annotated. Books on the Napoleonic War are annotated, "This is not the way it happened." One of the employees at Davidson in the 1800s was Marshal Ney, who worked under Napoleon, so he said this is not the way it happened, and he wrote it in the margin of the book. Well, so you read the book and it says "This is the way it is"; there's the marginalia and it says "No, it isn't". In a computer file, you just overwrite it and you can't tell. So this…. At this point I was living halfway across the state in Auburn, Massachusetts, just past Worcester, so I had an hour commute every day. And as I was commuting, I was thinking of how reading some stuff and writing some stuff and things being copied from one to the other, in/out. I got all excited about it and rushed in and before I got my coat off I was telling Len that I had had this interesting idea, pulling off my coat, drawing on the board, talking about subject at one level, and a higher security level object and a lower security level object, and information could flow and we have to stop it. How do we do it? I didn't have a good name for it so I just put up an asterisk-dash-property — *-property. So we talked about this and different approaches you could take to try to stop it. At some point we said if we don’t change the name of this now, we’re going to be stuck with it. But we couldn’t think of a good name for it, so we left it.

Yost: So in the 2005 retrospective paper, when you mention you put that up on the blackboard and you said we would be stuck with it [Ed. that is, the name], it was just you and Len in the room at the time?

Bell: Yes. Now, when we got ready two years later, we were writing a Multics interpretation and we looked at it and Lipner had never much liked it: *-property. And so we decided — lot of people had complained about it, and we thought about it and thought about it —even though it wasn’t perfectly descriptive we went into Lipner and said “Okay, so in the new version we’re going to call it the ‘No Write Down Property, NWD’”. And he said, and I quote, “I will defend to the death the use of *-property — with an asterisk,” which we wrote down and put on our wall. So he came around.

Yost: Did Steve Lipner ever convey that rationale of choosing you and Len LaPadula, having an engineer and a mathematician working together on this project? Was there the sense that that type of collaboration was the best opportunity to get a grip on this challenging problem?

Bell: No, he never said anything to me but he may have said something to other people. He may have said something to Len. At that first year, when Ed Lafferty, my department head, was running around trying to change my review, one of the things he said to me — or it must have been to us — was "How can I tell who's contributing what? You guys work as a team" and, in fact, Len had an interesting way of writing a report at MITRE. He had the old fashioned typewriter, you know, with the keys that strike, not the Selectric. When he got ready to do a report he got out a folder and a piece of paper. And he made a cover. You know, he labeled the folder and he put it in it. Then he'd do a table of contents and put that in it. And when we wrote together, we'd go to the folder, we'd pick it up, we'd see which section had not been written, and we'd just announce, "I'm going to work on section five" and we'd write it. We'd write by hand. Then later we'd go to the typewriter, type it out, and then whoever had written it handed it to the other person who would then edit it. So in some sense, it was awfully hard for us to tell who had drafted words because we'd go back and forth and it kind of blended. And Lafferty said "I can't tell if I put you on different projects if I get twice as much, or if one of you is doing all the work, or if each of you is doing it — can't tell." So, I think, from my point of view, we were a unit so they said, "You know, you need to do this kind of stuff. Len and David had been doing all this cooperation, you know, the book's very mathematical." In fact, we had a theorem and we had some debate about whether our theorem and our proof matched. It did. So I think they just picked us as a unit. He was the one who actually knew computers. The notion of having me do computers without somebody who understood computers would have been stupid. Maybe they had nothing else for me to do that year. Don't know.

Yost: As you were doing the work that resulted in volumes I and II, did you have any interaction with MIT and the people that had worked on Multics?

Bell: There was one meeting but I can't remember if it was during that time frame. But it wasn't much later where a bunch of the MITRE people went off for a meeting down at MIT with Multics people as well as some other MIT people; probably Air Force and who was it then? AT&T, or GE, or Honeywell? Some corporate person or group. And I can't even remember what it was all about. The main thing I remember was Peter Neumann saying, making reference to Dykstra as somebody many of them view as a god. I'd read Dykstra and didn't view him as a god, so I was a little surprised. But there was just that one time.

Yost: And at what point were you aware of the project that was being done? the research that was being done at Case Western?

Bell: That was during the same first year. And in fact, we went off to Case because they were having a presentation of some of the work they were doing and we went along there. There were two highlights there. One was I actually saw Clark Weissman in the front, the ADEPT-50 thing. And just as he was leaving — to go get an airplane — so he made some comments, but that's the first time I had laid eyes on him. And the other was that in their early modeling work, the first modeling paper, I thought they had made a math mistake. They had said that they could make use of a "preorder," which in mathematical relationships . . . . It doesn't have what's call "antisymmetry." A less than or equal to B, B less than or equal to A implies they're the same. Well, in a preorder, that doesn't have to be true. So they said, you know, "You can do this with a preorder." So I raised my hand, and I said "I think you need a total order." [Ed. Bell misspoke. He should have said "partial order."] And they said "Actually, no." So I sat back down. So they went along. At a certain place, they came to their theorem, and I stood back up and I said, "And here's where you need antisymmetry." So that made me feel good. So they fixed that. We read their initial report. We went to that meeting and that was all. Now, a little bit later when volume III was happening — was that where it was happening? Yes. In volume V [Ed. There was no volume V. Bell misspoke and should have said volume IV.] …. volume III, I was doing myself because Len had abandoned me. Some of their work got inserted into what I was doing but through the intermediary of Roger talking to Steve talking to me.

Yost: In what aspect?

Bell: The issue was if you have a hierarchy of files, kind of like Multics or UNIX, what is the classification level that you place on parent directories, compared to files? And both of us were making the same mistake — if it's important, it must be classified. But everybody was thinking that at the time. What they said was because you chase down a path from the root down to the bottom, you can only be going up. You read it, you go; you read it, you go; you read it, you go. What I said was, "But the parent has information about the lower one. You should….   The root is an exception, it has no parent. So from the root to the next one, you do this exceptional jump, and then you only go down." So I was going that way; they were going that way. I was ordered to use their way. I wasn't happy about it. So I called it "compatibility," I think I said, and I gave them credit or blame, whichever one you want. Now, we were both wrong. The fact is that protecting metadata, which is . . . the contents of a directory that tell you about a file is metadata, it's at a different level. Protecting the metadata, yes, it's an important thing, but that doesn't mean it has to be SECRET. Maybe what it has to be is "system protected", or some other category extension. The problem with their method is that someone reading a SECRET file; no, not what I meant to say. By altering a SECRET file that has an UNCLASSIFIED directory changes metadata in the UNCLASSIFIED directory, so by having a prearranged plan you can make changes in order on your SECRET file, or set of files. Someone who doesn't have a SECRET clearance can read the metadata at the UNCLASSIFIED level and get it out, you know, a kind of "covert channel" is what we call it these days. So no matter what, a successful engineering solution said that you had to deal with metadata differently than plain old data, regular user data, so you had this exception you have to deal with. Theirs set up this covert channeling situation. The workaround is you never go from directory to file as a step up; you go directory to a step-up directory — directory, file. So you have that workaround. So now you can't read the directory unless you're up there. That's the extent of my interaction with Case. And if they had a second volume I never got a copy, so I don't know what they did past the first volume.

Yost: So primarily, it was two relatively independent projects with most of the intermediary communication by Lipner or Schell?

Bell: Yes.

Yost: What were the greatest challenges in developing the *-property model, in the math behind it?

Bell: After we came up with the . . . After the problem crystallized, it wasn't very difficult. We were thinking about the sort of the engineering of it to figure out which solution we wanted but conceptually it wasn't hard. The two ends of the spectrum — there may be something in the middle — but the two ends of the spectrum are if you don't want information to go from the intelligence summary into the bowling scores, then you could monitor every single transfer. The other end of the spectrum is you avoid the situation of having both of them open at the same time. At the time, we were willing to say someone at the . . . . You know, this Top SECRET person, might have access to both of those. But not together. So you had your choice. We were going to insist that you ask for one at a time. So if you ask for the bowling scores, to write, we'd say "Sure." If you then asked for the Intelligence Summary to read, we'd say "Sorry. You're prevented." Or the reverse. You know, "Do you want to read the TOP SECRET Intelligence Summary?" "Yes." "Do you want to . . . . Now I want to write the bowling scores." "Sorry." So it was preventive and from an engineering point of view, what that said is if you can keep track of what's open and then know for sure that that's all you can touch, then you can prevent these flows just by which ones you allow to be open simultaneously. So the latter is what we chose. At a modeling conceptual level that was pretty simple. The complexities that came from engineering were in volume III, that I did after Len went off to NASA.

Yost: You later gave a talk “Looking Back,” and had a slide saying that development of a tool requires tinkering. Can you give some examples of tinkering in developing the Bell-LaPadula Model?

Bell: Yes. Actually, in that paper one of the things I did was purposely include the agonizing math symbology of the basic security theorem because the definition of state changed in every single one. In the first one, there were objects that were monolithic and so state represented …. included monolithic objects. In the second volume we had access modes, so the state had to include the ability to say which mode you had access in, which you didn’t have before. We also included the *-property, which in version I, security was only do you have it or not? In the second one, it was do you have it or not? and is it *-property secure? Does it satisfy the *-property? So you’ve added an additional definition of what security meant. In the third volume, people at MITRE were trying to make use of the modeling work we’d done in the first two volumes in at least two endeavors. One was to make a UNIX kernel enhancement, and they were going to try to secure it, for some version of “secure.” They would take the kernel and make some changes and additions to it. Enhance it. The other one was that a bunch of the MITRE people were involved in doing… helping with the Multics project, helping secure Multics. In both of those cases they started having trouble with using the model in the actual programming/engineering situation. The first was Lee Schiller was doing the secure UNIX enhancement, and he came up and he said “The *-property’s driving me crazy. I’ve got this scheduler,” he said. I had a vague notion of what a scheduler was but as he talked I understood it better. “And its function in life is when a process finishes is to store everything away, all of its state, all of its data, figure out who’s to run next, and pull it off store, and then kick it off. If I am switching between UNCLASSIFIED and SECRET, I have to read and write UNCLASSIFIED, and read and write SECRET. What level am I supposed to make my scheduler in order to be able to do that? I can’t do it.” There was nothing I could say. That highlighted the fact that in the earlier versions we were not making distinctions between types of subjects, which turn into processes or jobs or something within the computer system. They’re all the same. In an actual computer system, they’re not. Some of them are part of the kernel, part of a lowdown . . . part of the system, and part of them are just plain programs, and text editors, and things like that. So in order to do that, what we had to do was make the distinction between subjects that want to be free of analysis and care before they’re installed and the ones that want to be a scheduler or something similar, that can open things up simultaneously but we promise not to do the kind of bad transfer. We didn’t make it absolutely clear in the paper but it was in the air that things like the scheduler, if they are going to be accepted to not transfer, you have to make the argument. What’s the argument? That the scheduler doesn’t transfer the information, doesn’t do something wrong, throw in a little TOP SECRET file into the bowling program. So we ended up calling the subjects that didn’t have to be limited by *-property, called them “trusted.” And the other ones were “untrusted.” We recast the *-property to say “if you’re not . . . if you haven’t been labeled, if you haven’t been approved as a trusted subject, here are your limits.” Now what some people did was that they ignored the step of careful examination of your trusted subject program and they started saying if this is too hard we'll just bless it and call it trusted and therefore, everything's okay. Well, that was not the intent. That's not was it was supposed to be. So there's trusted subject.

Another one was, Lee also came up and said, "The way you have described 'trusted subject', checking for opening up an object for *-property says check that security level against everything that's open. So if I have 100 open, I've got to make 100 comparisons. For my programs, they always have a workspace that I read and write. So what I'm going to do is just compare that new security level to that one level. That's one comparison instead of one to 100." I said "Good idea." So I went back and put that in because one of the things in those days is that we were feeling our way. People were comfortable if they could say "I did exactly what's in the model." They were less comfortable if they said "Well, I took what's in the model but I made these little changes." It was always more comfortable if we found something we needed to do, to pull it back into the model and look at it in a purely conceptual way and were not having to worry about registers and compilers. And then they could say "Yes, it's just like that." And this came up in Multics as well.

And the other one was providing a "hierarchy," a file structure that included what we talked about, the Case Western and the compatibility — what's the order of the elements? "So you want to make a new object? Yes? Okay, has to be at or above its parent." So that kind of state and that kind of check wasn't in the earlier version because it wasn't part of the definition of state. So as we proceeded, we kept having to add those things. Now after we finished the first three models I went off and did something useful. Len was helping our space program. And the year following, when we ended up doing the Multics [Ed. interpretation] we found a lot of other things that needed to be changed, too, because it was specific to Multics.

Yost: Can you go through those changes that were done specific to Multics? At least the main ones.

Bell: Yes. Well, at that time, the project leader was Ed Burke. And Ed Burke came to us and said “We’ve got two problems here. The people working on Multics, they see your reports, they got the three reports, and the leap from these reports over to the way Multics works is a big leap. You know, these are software guys. So that’s one problem. The other problem is the one I just mentioned; the model keeps changing every time you put out a new edition. So it’s hard for them to track it all so we’d like you to get it all in one place and then help us bridge that gap to Multics, make it look more like Multics.” So what we ended up doing, by volume II we had organized the rules about changes of state into things we called “rules.” If you would like to ask for access in a read mode, this rule applies. We kind of constructed it conceptually so any set of rules didn’t step on each other’s toes — that is to say, for any particular input there’s only one rule that’s going to try to tell you what to do. So any set of nonconflicting rules constituted a system and we had 11 [rules], because it was everything we could think of — 11 of them. For Multics what we ended up doing was deciding that we needed to make some different rules that were more specific to the way Multics did things.

One of the things we had left out in the first three volumes was the general treatment of control that an individual subject has about changing attributes. That was a conscious decision because the computer systems we were familiar with . . . there's just no homogeneity. All did it in a somewhat different fashion. So what all we said was "There's this control thing that tells you what you can do." Multics had a very specific way of deciding what you could and couldn't do. You can write in the directory, you can do it. End of story. So our rules had this general "control" system. Okay, well, in Multics it was different. There also were things that were in Multics that we hadn't addressed. "Please [dismount] this file system." Well, it actually was possible to look at the rules and say "Use this rule recursively." You can do it. We ended up writing a rule for that but at this moment I am not positive that's in. I came back to Multics when I was in NSA, so I don't remember if that recursive rule was put in in the Multics interpretation or later. But that's the kind of thing one wanted to be able to do, dismount and mount a file system. We had the basic tools. We didn't have a rule that said "It's just like this." So what we started with were eleven rules. We wrote an additional seven or eight. Didn't obliterate the first eleven; we just added extra ones. So the model for…. It's kind of parameterized model. Which rules would you like? Those? Okay, as long as they don't overlap, you've got one. For Multics we said, "Okay, we need new rules." So we wrote the new rules for the kind of things I just mentioned. Now, one of the…. Neither I nor Len knew Multics. So we went to the software engineer types and said "We need to arrange time to talk to you about Multics." And they said "We don't have any time. Take this book," and they handed us Elliott Organick's book, *The Structure of Multics*, which we didn't think was helpful. But we went off to read it, and it was hard for me because I'd never taken operating systems theory or anything like that. But that was our basis for doing it and when people reviewed our report they told us things that had changed. I mean, Elliott Organick wrote his book while the system was still being developed and it had changed. So we relied on the book and we weren't completely up to date. So that's what we did.

Yost: Were there…. Because security was a consideration from the start with Multics, were there things that made Multics easier or was it actually more difficult because to truly have a high assurance system you needed to go take a different path?

Bell: Well, as a matter of fact, this was the first time we tried to deal with a specific system. Early on, we had come up sort of with a three-pronged attack. First, you needed to be able to describe what security meant. They you needed to try to get yourself some general results, analytical tools. And that would be all fine and good, but if you really wanted to field something you needed to show that you could use it on a real system. Multics was the first one so I can't compare Multics to something previous. I *can* compare later to Multics. And the fact that it was built with protection — security — as one of its aspects meant that unlike some systems that I ran into later, they hadn't already put themselves in a hole somewhere. Whenever I was off on a project where somebody was trying to achieve security or understand what they're . . . . You know the first thing I would — you know, as a joke — I'd say "Where's the Organick on this operating system?" and if I was not understanding something I would say "Compare that to Multics so I can understand it." And Multics became, you know, was the prototype for B2 under the TCSEC [Ed. *Trusted Computing System Evaluation Criteria*].

Yost: At the time you were developing this model, did you have a sense of the importance of what you were doing and the impact it would have?

Bell: Ah, no. Security, in my experience — you know, one of the blind men — it was a niche activity. Nobody cared about it. We gathered together. We told everybody it was important. And everybody ignored us. And gradually, the support for doing this general stuff and the prototypes and all of that sort of stuff started falling away. It was only later…. You know, we figured that we were doing something that nobody would ever care about. It's easy to look back and say it's a good thing we did that stuff. But at the time, you know, the project managers and all those people, even the people in the Air Force, they were having to fight — *struggle* — to have the money and you know, keep the money and keep things going.

Yost: What was the level of dissemination of volumes I through III? Who saw those reports? There was an emerging but rather small computer security community with the DoD, MITRE, SDC, NSA, and RAND, and other places.

Bell: The way things worked at MITRE . . . . I largely didn't know. I didn't know the answer. When we prepared to publish a report, we would fill out a form and send it to the publication office. We had to give them the charge code so they knew who to charge for the publication. We had to give them an abstract and the title. And in the fullness of time they would send us a number back. Our report had a distribution list that we got from our project leader. And there were some external people, but mostly it was Hanscom Air Force Base Electronic Systems Division. Most of the time, all of those reports were then put in for public release and when public release came through, and when they were issued as an ESD document, they were then made available to anybody who wanted to request them from DTIC, I think it's called — the Defense Technical Information Center — that sounds about right. The people that got it that way, I have no idea. I have no idea. What sometimes happened was that later I would get a letter or a note from somebody someplace else out of the country asking for a copy of the report. Well, when I got a letter from an academic in Poland in 1973, 1974, I passed it immediately to the publication people because it wasn't … I wasn't supposed to have contact with people like that so I didn't. Now, one of the things — funny story — I told you that to get a report published you had to do those things. And for volumes I, II, and III we had done that. For the fourth volume, it ended up over 100 pages, quite long. And so we sent it in and we got back later, a folder, envelope from the publication office and a note that said too many people had signed up to receive it. See, after you sent it in, one of the things that was shipped out to all the employees at MITRE were these — we called them the "Rainbow Sheets," — and the yellow ones had notices of reports that were coming out. So if you saw something you loved, you'd just mark it down, you want that one, you'd just fold it over and send it back to publication and they just put it on the list. They just added it to the list. So we got this list and they said "You have too many people. You have to decide. You have to mark so many people off." This never happened to us before and we didn't know what to do.

Now looking back, the right answer was . . . . The reason was our project had a budget and within the budget there was publication costs and publication of the paper depended on how many copies and how many pages. And we had blown our budget. Now, they didn't tell us "You blew your budget". They said "Mark people off." If we'd been clever, we'd have gone to project and said "Can we ante up some more money and let everybody have a copy?" Well, we didn't think of that so we had to mark through, and so there were some people that signed up for it that didn't get one because we had blown our budget because it was a long paper and the list they sent us was over 100 people at MITRE. At the time, MITRE was maybe 2,000 or 2,500 people. Over 100 people. And we were kind of surprised because we only knew of about 20 people that were working on security, as far as we knew.

Yost: And to what do attribute this broader interest? Was there a sense that here's something that's really important and influential?

Bell: Well, what we took from it was that we had narrow vision. You know, we thought the only people that cared about this were the people we knew working on it. And that surprisingly, people outside our group were paying attention. So it was more just a matter of "Well, that's surprising. We could actually start a club because we have enough people that, you know, it's not just a dozen of our closest friends."

Yost: Were there readers of volumes I, II and III outside of the core group you were dealing with that gave feedback that was useful in bringing everything together in full?

Bell: No. Most of the work that Len and I did on this, we thought we were going to be talking to the Multics developers, but they just gave us the Organick book and said don't bother us. So we mostly, we did the work just in our office. Nobody interacted with us.

Yost: What about in the mid-1970s? Were there any strong critics and what was the nature of their critique?

Bell: The critics I became aware of were in the late 1970s, early 1980s. There were critics earlier, but I didn't know about it until the 1990s or 2000s. Some of the critics said there's nothing there, it's unimportant. And what I have always said is it's not the theory of relativity but it's not worthless. It's in the middle there somewhere. You know, it's a useful thing. I worked on security from 1971 to 1974. I then spent from 1974 to 1983 doing things besides security. So what is that? Nine years. So in 13 years, I had spent four doing security. I didn't think of myself as a security guy. But in 1983 I was deciding whether I wanted to stay at MITRE for the rest of my life or if I wanted to look into other things. And Steve Lipner called me up and asked me if I would be on a panel at the Oakland Conference, which I'd not heard of before, on the value or usefulness of Bell-LaPadula's security model. Or maybe it was models in general. And I said, "You know, I've been out of the field so long I'm not sure I want to go there." And he said . . .

positive things and I said “I just don’t want anybody to cut my knees out from under me.” And he said “The only person capable of doing that in a public arena is Roger Schell, and he’s on your side.” So I agreed to go, but I got hit with the criticism, which is . . . .   The criticism was this doesn’t give us a definition of security and all it shows us is that this weak definition is inductive. But I had prepared a trap and so I pointed out with examples that there were good properties that were inductive and not inductive, and bad properties that were inductive and not inductive. And knowing that you’ve got an inductive property is a good thing to know because inductive simplifies your life. So that . . . I mean, that’s one of the critiques. There was a little bit later critique that said that . . . .

I don’t like model wars. Yes. People started saying “It’s too specific — it’s just the Department of Defense.” People said “It’s not abstract enough.” People said “The definition, as it stands, is wrong.” Now, Len and I set it up . . . . Here’s what we’d say security is. Three things, actually four with compatibility, you know, the hierarchy. And some of peoples’ criticisms were, “I thought of something else that you haven’t included.” My answer was “I wasn’t addressing that.” So, there were modeling wars, but they weren’t very pleasant to go through.

Yost: And was John McLean’s critique part of that Oakland event or was that later?

Bell: Yes, it was. [Ed. Bell misunderstood the question. John McLean was not part of the 1983 panel. His criticism was at a later Oakland conference.]

Yost: Now, Ken Biba was also at MITRE, correct?

Bell: Yes.

Yost: And in developing his integrity model, did you have any interaction with him?

Bell: No. He came a little bit after I did and he had a master's degree — I think it was Case, I'm not positive — where he had done one of the early covert channels. And I remember seeing the draft or the final of his report, but I hadn't interacted with him before that.

Yost: And can you comment briefly on his model and the Clark-Wilson model and what you feel they do or do not accomplish?

Bell: Well, I mean, part of my reaction to the Biba model was that part of it was just an elaboration of things that Len and I had mentioned. We had talked about integrity at the level of "we have to control what people can write." We also talked about security and sabotage, and what people now call covert channels, but not in great detail. We sort of mentioned them, that these are things to worry about. I think the . . . good importance to Ken's work is to get people to think more about the limits on making changes to things. Some people took his strong integrity, if I'm getting the terminology right, and applied it with the same set of labels as confidentiality and they got trivial systems. Well, in some later work I did as well as things that were done in GEMSOS, and the Multics interpretation — not my report, but what was done at NSA — showed that the notion of using the same labels we're used to for classified documents for every security thing you could think of is shortsighted and wrong. That if that's appropriate, you should use it but you really need to think about whether to use it. In GEMSOS and in my later work on lattices, the fact is that…. And a Biba integrity lattice and a confidentiality lattice can be embedded in a single Boolean lattice. And a Boolean lattice is just like the DoD lattice. So, by an appropriate set of personality modules, you can take a DoD lattice and do both simultaneously, all in one place, which is as I understand it, what GEMSOS did in the first place. They didn't actually separate their integrity from the confidentiality; it was one monster lattice. So. Now who else did you ask about?

Yost: And the Clark–Wilson model?

Bell: Clark-Wilson was very interesting. It was . . . . Since it derived from a set of rules that were prevalent more in business, in talking about integrity, it got a lot of people interested and excited. A lot of people actually decided that there was actually different kinds of policies. Here's a worked example of an non-DoD, non-lattice policy. So it created a lot of excitement, a number of papers. The following year three or four Clark-Wilson response papers came out in Oakland. Then, or not long afterwards, there was a workshop, whose name I can't remember but the initials I can. Up at a university in Boston, whose name I am now forgetting — Brandeis, maybe — it was called WIPCIS and we thought it was like heavy metal, Whip Kiss! — to talk about Clark-Wilson and the extensions into the business world. Also about a year or so after that was the Chinese

Wall paper by . . . Oh gosh, who was that? Brewer and Nash, I think it is. What they did was to come up with a security model that reflected a U.K. financial rule about financial analysts, reporters, and the . . . . If we go back, let's see. Clark-Wilson. The most interesting thing from my conceptual modeling point of view is that instead of having duples — subject and object — they have triples: subject, the program you're running, and object. So that was different. That was new. The Chinese Wall said that if . . . you know, suppose I'm a financial analyst and I'm doing oil, so I look at the oil stuff. But if I decide I wanted to specialize in Exxon Mobil, I can choose to do that and I can then get more detailed information of Exxon Mobil. But that discretionary action of choosing to become a specialist in Exxon Mobil prevents me from choosing to become an expert in BP. So, by a discretionary action, you've blocked yourself out from other directions. Well, this was . . . . It came from a totally different field and it suggested a kind of action that just hadn't been dealt with in the DoD. The notion of saying "I think I'll be a photo interpreter." No, you don't get to do that in the DoD. Somebody chooses and off you go to do it. Well this was . . . that was the newness and importance there and part of what became clear to me a little bit later writing lattice papers was that we had all of these papers popping up saying "I've got a new kind of discretionary access control. I've got a new commercial policy here. I've got a new financial policy here."

There are two cases. If they really were different, then the vendors — then the people who were building the computers — had a hard problem. They had to choose which ones they were going to support and if they were different then they were only going to cover a couple because they looked to be having to build their software differently for each one.

If they only appeared to be different but they were similar then they could write the base software the same and have personality modules. So if you want to fake Clark-Wilson, you could fake Clark-Wilson. If you want to fake Chinese Wall, you could fake Chinese Wall, and so on. As it happens, that fed into something that I had undertaken when — let's see, it must have the Computer Security Act of 1987, that sounds right — new set of requirements were put on NBS [Ed. National Bureau of Standards], which is now NIST [Ed. National Institute of Standards and Technology]. Which . . . . They had a bunch of computer security requirements they didn't have before and one of the things they did was to hire a Trusted Information Systems to give them some lectures and tutorials on things. So one of the things I got tasked with was to try to explain how a lattice could be used in a commercial area. So I put together a little presentation about information within a company. You know, there's financial stuff, there's strategic plans, and there's engineering. And some things could be strategic and financial and some things could be engineering. So we explained all this and I had the slides and put it up. And so they listened and asked a few questions, and as I was sitting down, Marv Schaefer, he was the Chief Scientist at the Defense Computer Security Center when it started, and he was at Trusted Information Systems later. He said, "Of course, we don't know how to do lattices with 'ORs.' " And as I was sitting down, I thought to myself, on the other hand, "Why can't we?"

I had been working to understand why we couldn't do "ORs" and eventually, I figured it out. The reason was that in lattice theory for boolean lattices, a boolean lattice is determined by things that are just above the empty set, which are called "atoms." And what we had done in the security business was to take our security levels and call them atoms. And what that meant was that we could only do “ANDs.” So you take alpha AND beta and alpha AND beta AND gamma. But if you want to do an entire boolean lattice where you can do ANDs, ORs, and NOTs that — what I call the “policy alphabet” — wasn’t at the bottom, It was up here in the middle. And in fact, in lattice theory, there’s a way that says if you care about these, this is exactly what you need for your atoms. So I wrote up a paper called “Lattices, Policies and Implementations” where I said this is the way lattices are, particularly boolean lattices, and the biggest and most important thing out of that is that there are many different isomorphic ways to define a boolean lattice and one of them looks like the DoD lattice policy. And one of them is built by symbols — uninterpreted symbols — that are combined by AND, OR, and NOT, which is almost every policy you can think of. So in math, lattices are isomorphic — and they’re all the same — if it’s one of these versions, it’s all of these versions. If you have a policy that looks like one lattice version then it *could* support any of them, if you have a personality module. And if you implement one of them, and you put a personality module, you can do them all. Supply them. I said that one year. And the following year . . . . I then decided that what I was going to do was to check everything in the literature that said “This is not a lattice model. This is not a lattice policy,” and see. And I gave a constructive answer that all of them *are* boolean lattices. Now, some of them are boolean lattices with a couple of other things. Like Clark-Wilson’s limits include a few others. Some of Clark-Wilson includes more stuff. But most of Clark-Wilson is a boolean lattice. If you can express it with AND, OR, and NOT — boolean lattice. So I wrote that one up. It was interesting that the “Lattices, Policies and Implementations,” the first year

I submitted it, it wasn’t quite finished. And they said “This is not finished” and they didn’t accept it. The next year, I submitted it and one of the reviewers said “Everybody knows this.” So I gave up on Oakland and sent it to the National Information — what did we call it? — the National Computer Security Conference. And at that conference, large numbers of people said, including the guy at DEC — same name as the transistor guy [Ed. Bill Shockley] — said we were doing this at DEC and we were going to view it as a trade secret. A lot of other people said “No, we didn’t know this.”

So after that conference I went to Jim Anderson and said, “Did you see my talk?” He said no. So I briefly described it and he said, “If you could address ORCON — originator control — I know people who could use it tomorrow.” So the following year I wrote a paper saying there are all these things: multinational sharing, ORCON, Clark-Wilson — I don’t remember all that I used — Chinese Wall, you can do them all and this is how to do these things. You can do these things. So what that did was answer the question are these different and now we’re in trouble? Or will personality modules do it? The answer is we can do it with personality modules as long as you have a Boolean lattice and AND, OR, and NOT.

Yost: Speaking of Jim Anderson, did he have a continuing interest and were you in contact with him as you worked with Len on the development of the model? You obviously had some contact with Steve Lipner and Roger Schell, but any contact with Anderson in your first half decade at MITRE?

Bell: I didn’t meet Jim Anderson until I started going to the Oakland conference in 1983. I went for seven or eight years straight, then once or twice after that. I met him there and we never worked together but in computer security conferences, people would argue, argue, then go to dinner. So there were various groups that went off to dinner and Jim and I had a good relationship. He recently died. Just before that, turned out he was working where my wife works so she worked with him as well. And we invited him to come and stay with us, but he was off to see his son and grandson so he stayed there instead. While we were doing the modeling I did not know him at all.

Yost: Okay. Do you want to take a short break?

Bell: Sure.

Yost: You mentioned the Oakland Conference got you back into the computer security research area, and that was 1983?

Bell: Yes.

Yost: Is that also the year that you left MITRE to go to NSA?

Bell: Yes, actually it was . . . .   I was . . . . Let’s see, the Oakland Conference that year I think was in May. I went to work at NSA in July and so I think I must have already contacted them. As I said, at my 10-year anniversary I started thinking to myself, "If you're going to stay here forever, decide. Don't just slide into it." I looked around at various things and I ran into Marv Schaefer at the MITRE cafeteria. He had already become Chief Scientist there and Roger Schell was the associate director, as a full colonel. Marv saw me. He was up talking to the MITRE people about security. I had figured security — lost corner. You know, I was off doing other things. I was putting systems in the field to fight the war. So he saw me and he got all excited about it. Eventually, he and Roger convinced me to come down to talk to them about working at NSA. And I think that happened the same time that Lipner said "How would you like to come [Ed. to Oakland]?" So that was . . . . So going off to talk there was sort of as part of a tumult of deciding to go off to NSA. They invited me to come down and — my ex-wife reminded me . . . . But she came down, so she was going to see what kind of job she could get. And I went off to talk to them to talk about what was going on. We did some house hunting. And Marv and Roger suggested I go over to some sort of picnic or something that was going on. So we went driving off into the hinterlands of Maryland, and pulled up and went to a barbecue at Steve Walker's house.

I didn't know who Steve Walker was at all and as we drove away in our rental car, before we drove home — I forgot but my ex-wife told me — [I said] "Someday I want to work for him." So, went to NSA, then went to TIS with Steve Walker. What Marv and Steve [Ed. David misspoke. It was Marv and Roger] . . . . There was a program at NSA — I think it was COINS — and they were going . . . . Let's see, what was going to happen? They were either re-implementing, or they were merging it, or they were making it a high-level A1 system and they needed somebody. I'd been doing program . . . . I wasn't a program manager but . . . . the Air Force program manager. I ran the MITRE project that supported the program manager doing acquisitions. So I was [unintelligible] [focused on practical systems-in-the-field] because [I worried that] the geriatric Soviet Union could do something stupid. So I was built for this [project]. So when they contacted me I said I really don't want to do "research-y" things. I want to build systems and put them in the field. And so they said what we've got is this COINS-like system and it needs to be really high security. So I thought about what the job was going to be, and I took this long time trying to figure our this benefits package, which was not easy, comparing a private nonprofit benefit package to the government benefit package. And in the 1980s . . . . You know, the 1970s were this amazing inflation time so one had to worry about, you know, what's the prospect for not losing to inflation here. So I did all this calculation. They had sent me the package, you know, to sign on the dotted line you're going to accept this job. I finally figured out that I could afford to go for two years, after which I'd have to decide. Either I would get something to increase my salary or I was having so much fun that I would do it anyway. So I called Roger and said, "Okay, I've decided I'm going to do that." He said, "Oh, that's great; that's wonderful. We have a different project for you."

Hold on. Hold on! So he said, "Our new idea is for you to be the Deputy Director of our R&D group." I said, "R&D? That's what I told you I didn't want to do." And they said "Well, I think this will be a really good fit. I think you should go talk to George Jelen. He's currently at the Harvard," No, what's it called? The government place at Harvard, the government institute, the Kennedy Institute, or something like that. "Anyway, he's up there so maybe you should get together and see how it works." And so we met at Grendl's, something like that, somewhere in Harvard Square — a little hangout. He was pretty much a blue-suit kind of guy, and I came down, we talked, and I couldn't imagine two people being more different. But he pointed out that his plan . . . . What he wanted to be able to do with his position running the research group was to be outward focused — help grow it, keep it going And what he wanted me to do was deal with the inside — build the program and all that kind of thing. So I hemmed and hawed and decided that would be good. As I said when I was retiring, I've never had a boss I didn't learn something from. George Jelen said something to me then, or later, that was one of those. Which was, "Nobody should ever take a job they're certain they can do." It struck me as strange at the time, but I think it's a good thought. You're stagnant if you do that. So I went down there to do that. It was kind of interesting.

I went to Friedman Auditorium with everybody else that came in that day; hundreds of people sitting there. So they stood up and they said "Welcome everyone, we're getting ready to do orientation. Would the following two people come up." And I was of them. And they ushered us off to the side. And I was at work at 11 o'clock. They didn't send me through orientation. I didn't get to take the basic cryptography class; they sent me off. So I got to the DoDCSC/NCSC, which was at that time 60 or 70 people at that time, jammed together because we didn't have enough space as yet. So I came in and they introduced me, and explained the badge. I'd already done the polygraph test in the morning, which they did with kids' gloves, and I had close to a month of a honeymoon. I was introducing myself, I was wandering around, learning what to do. And then there was a crisis and they called me in and said “We need you to work on this project called BLACKER Phase 1.” And I said “Okay, what is it about?” They said “It’s a network system that has a lot of security in it,” and I said “I don’t know anything about networks.” They said “Well, we have got an answer for that.” [Laughs.] “Ray McFarland and Dave Solo are going to get you spun up to speed.” So I spent one morning with them and I came out with my head spinning, not having absorbed all that much. But I ended up the lead for . . . . Remember at MITRE I told you about soft shell?

Yost: Yes.

Bell: Well, what we had was people from three different organizations at the Center supporting the project being run out of a different group, the R Group. And none of them reported to me but I supposed to coordinate them. [Laughs.] I had no hammer and I had very few . . . . And so off I went and for the next 18 months it was all BLACKER, all the time. What we were doing there was building a networking system that originally had [fourteen] different components, but ended up getting narrowed to three that was trying to do A1, in a network, geographically distant, three different kinds of devices, and nobody knew what A1 meant for this.

Yost: And that was certainly the first attempt to build an A1 network, correct?

Bell: There was another one. I think that SACDIN was also a network kind of system. I don’t know the details of SACDIN well. But this one, if you think of it as routers that have on the back end, stuff to do really high level encryption and a server somewhere on the network that supplies you keys for people you're allowed to talk to. That's what it was. Now one thing we skipped over was *TCSEC*. When I first got there they hadn't quite published The Orange Book. They had it in draft form, but they hadn't published it. So they gave me a copy, turned out not long before they published. Said "Take a look at it." So I'm looking at it. Went to the glossary. *-property — had it wrong. So I said, "You've got it wrong." They didn't change it. [Laughs.]

So back on BLACKER, what we ended up having to do . . . . It turned out the government and the contractor, both wanted contract changes. Well, this is the best of all possible worlds because you get something, he gets something, so we compromise. One of the things that they wanted to fix, at the very last minute during the — not best and final — but the final negotiations before they sign it, when there were no techies around, the government inserted some requirements on this contract that said BLACKER Phase 1 will be "designed in accordance with A1." That's a very bad contractual statement. So what the government wanted was some clarity on that. Actually, the other side wanted something too, because it was hard to know what that meant. People who understood it, like Clark Weissman, you know was there . . . was their senior security guy there. And they were going to be doing formal verification. They needed to know what they were actually responsible for doing. So one of the exercises we did was that we sat down — Government . . . . C3. We had C3, we had C2, another C3 division, and I don't think we had a C5 yet. But what we did was we sat down. We had a group of seventeen. Each of us had an Orange Book. In our initial thinking what we said was "Each of these devices needs to be good and secure. Think A1, maybe." We had three devices. But on the other hand, that's not enough. The whole system has to fit together in such a way that we can say to ourselves "Boy, that looks like A1." And if you think about a router and a place that throws keys at you, the third device was something to generate a key. And these were symmetric keys so both sides had to get the same key. So the general notion was that we . . . started through A1. And we would read it a phrase at a time and we would have to decide "Does this apply to the whole system, yes or no? Does it apply to the front end? Does it apply to the access control center? Does it apply to the key development center?" Phrase by phrase through all of A1. And having done that, we then wrote an annex to stick into the revised contract that said "This is what the contractor is required to do." This ended up being a subset of what The Orange Book says because some of the jobs, things that you needed to do, were clearly government jobs. Some of them were part of them and part of us. But more than that, some of the jobs on the government side, it was unclear whether this organization or that one was going to do it. Is this C3 or C2? So we went through it and we pulled out those parts. We put it into nice contractual language, Annex O, and that became part of the contract. So that became our new definition of what does it mean to be, for the contract to meet its security requirements was Annex O . It wasn't The Orange Book. Now, Annex O was parts of The Orange Book, explaining who had to do what.

During that same set of meetings that we kept going to Los Angeles for, the Celtics were playing the Lakers in the NBS finals, so we sat and watched. And because it was, because we were on the West Coast, after it was over the band came on. So I was listening to the band and I’m thinking to myself, we’ve got all these requirements but what are we going to do about the model? How? This system doesn’t look like Multics. So I’m sitting there listening to the music and suddenly I thought to myself — remember the routers? — the basic notion was the router with this little back end hooked onto it. We’d say “Okay, David wants to talk to Jeff.” I’d say “Jeff.” Back end would say “Oh, you think you’re going to talk to Jeff. Well, what that really means ‘Yost.’ What level are you talking at? Looks like SECRET. Do I have a Yost SECRET key?” If the answer is yes, bingo! If the answer’s no, then it’d say “Well, let me check.” Then it’d say “I’d like a key for Yost.” If the answer was “yes,” I got the key but so did you. So the keys went both directions. If everything was up and running and there was this minor hiccup, then the magic happened. If we were turned down, then you know you’d get the message “What are you doing, you idiot. You can’t do that.” So that’s what it was doing. So what I ended up thinking was we have no subjects, we have no files. Then I thought to myself “Wait a minute. What the ACC is doing is deciding — mediating — [is] whether you can connect.” So the objects are these connections. So we have a SECRET connection between us. We have a TOP SECRET connection between us. We have a CONFIDENTIAL connection between us.

Those are my objects. It didn’t look like any other object I’d ever seen. And who am I? Well, I’m not a process. It’s a host, a distant end, an IP address. So now this is from the encryption business, what they were calling then a dual catenet [Ed. catenated network]. So, as I was talking . . . . I was saying “Jeff,” and I was saying it in a nice RED form, classified. The front end would say “Jeff,” “Yost” at BLACK. So you start off RED, it would figure out what the appropriate BLACK address would be, [and] it would encrypt all of the payload. So now you have a nice encrypted load that goes over here. The other end says “Okay, seems to come from this person, now do I have a key? Ah, yes.” So it would come in nice and BLACK, it would decrypt it, become RED again. Doing all this good stuff. So the pipes are by level and it’s host to host, just to the end IP address.

So the music’s playing. People are dancing. I’m searching for paper and writing stuff down. So one of the requirements that the contractor had was showing that they’ve got a security model and it all works. Well, it ended up, I wrote it. So this was a case where you had the modeling tool kit and looked around and said “Do we have anything that looks like connections and hosts?” The answer was “No, not exactly.” No. So I wrote an interpretation — “A Network Interpretation” was the paper, I called the paper — which sort of laid out if your subjects are these and your objects are these then you end up with slightly different rules — an object is defined by two subjects at a level. Never seen anything like that before. And these objects were both <u>read</u> and <u>write</u>, <u>view</u> and <u>alter</u>. So they [Ed. the contractors] knew that they had to do that at some point but I just thought that was stupid so I wrote the paper, and then they could just reference it. And that’s what we did. There were lots of ground-breaking things in that project.

Yost: In addition to the application with BLACKER, what was the impact or influence of that model? Do you know of other network systems that drew on that model?

Bell: Well, one of the things I did in the modeling was to . . . I think that was the place where I first talked about trusted subjects with limited range. In the original version of doing the schedulers what you say is you have to be wide open. And I think that was the place where I first said what we're going to do is you're trusted between this low level and this high level. Above and below it, you're not trusted. And what that did was collapse the entire notion of trusted and untrusted. Untrusted said those two are the same so you're never trusted. The scheduler is trusted is top to bottom. But you could also say I've got a sanitizer that is good from SECRET to CONFIDENTIAL, so it's allowed to make a copy from SECRET to CONFIDENTIAL but CONFIDENTIAL to UNCLASSIFIED is not.

Now that came from sort of a network view of the world. But later when I was at TIS and we were trying to figure out a way to make the Mach operating system — M-A-C-H — trusted I said "You now, if you're going to do trusted subjects you might as well do it that way." And so then they did. Now, with regard to networks, I don't know of any other place that people used it. Who was it? VSLAN Forgot the name of the company that had a B2 networking system. [Ed. Verdix] They didn't use that. And Boeing didn't. They did their own work. I might be wrong about VSLAN. I don't remember about what modeling that they used.

Yost: As I understand it, you also authored the model interpretation of SCOMP. Is that correct?

Bell: What happened . . . . While I was at the center there was one organization that did product evaluations. That particular organization was run by Mario Tinto at that time and during the SCOMP and during the Multics, people that were involved came to me and said “Do you have time to review something for us?” My memory is not as good as I’d like on this. I think it was, in both cases they asked me to take a look at the modeling and the “modeling interpretation,” is what I called it, from The Orange Book. The Orange Book says there will be a model and you’ll show that it applies to and matches your system. What frequently happened when people were coming for evaluation, they didn’t understand the modeling stuff very well. They would say that is our model over here. That is our system over here. And they didn’t write anything to hook them together. So maybe it was the SCOMP. Mario said would you take a look at this? I said “Sure, I’ll take a look at that.” And I came back and said “The bad news is they don’t have a connection between what they’re doing and the model. And the good news is I volunteer.”

This is another case . . . . One of what I realized in the early modeling that bouncing back and forth between the engineers and the modelers helped everybody. Another case. I put together this . . . . The SCOMP paper had said here’s the stuff at the interface, the kernel. So I went through all of that and I said “For each of those functions, here’s how you represent it in a version of the model.” Remember, you pick out the right rules so that it fits. Say this is how it works. So I put that together, handed it to Mario. Mario handed it to them. They said “Oh, so that’s what you’re doing. Well, in that case, three, seven, and fourteen are not visible at that level; so it should be dropped out. But we’ve got three new ones that should've been at this level and we think it works like this." So we had a couple of iterations back and forth deciding which stuff got submerged, which was actually visible, and where I didn't . . . . In some cases I thought something worked one way and I was wrong and so they'd say what was wrong, so we worked until we got it figured out. So we finally got that report so that they had a version of the model and a way of hooking it over to their DTLS [Ed. Descriptive Top-Level Specification] or their FTLS [Ed. Formal Top-Level Specification], their specifications.

Yost: Are there any other projects that you worked on at NSA in that time before you left for TIS?

Bell: Nothing technical. I spent . . . . Well, that's not completely true. Most of the time I was being a manager as well as a technical lead. There was a brief period of time when Bob Brotzman, Dr. Brotzman, who was head of the center at the time, asked me to go around the center and figure out what we might need to do to change how things worked. He sent a memo around that said that I was to be his "alter ego" in trying to look at operations. So I made up note paper saying "C/AE" — alter ego. Unfortunately, it made everybody defensive and nervous. "Here's the director's guy coming in to criticize what we're doing." So I went in to talk to people about what they were doing and have a conversation about what we might want to think about. Everything, I said, is kind of a possibility. They immediately looked at me and said, "Can we do it now?" I went [implied gesture followed by hand clap]. And one of the things I said was, "You know, it's really hard for people to understand what's in The Orange Book. It was written in a rush. The glossary was not written by the old hands. It was written by the newer people because they concluded late in the game they needed a glossary. And there was a press to get it released in time to go to a conference, or I forgot what the press was. Parts of it are just really hard to understand. I mean it's hard for me to understand."

And what I always did when I couldn't understand it, I'd go in and talk to Roger and say, "Will you explain this to me?" Then I'd go to Marv and say, "Please you explain this to me?" Then I'd go to Dan Edwards and say "Can you please explain?" I got three different answers and in the middle somewhere was what does this really mean. So one of things that I had been thinking is we really need a "Children's Guide" here — a "Children's Guide to The Orange Book." Well, the people in charge of writing documents, that's where the Rainbow Series came from, said "Oh, we can do that. We've been planning to do that." So this proliferation of "All About Mandatory Access Control", "All About Discretionary Access" — that's all my fault. What I needed was the simple version of what [all] this means and what we got was deeper versions of what individual phrases, individual terms mean. So other than the managerial kinds of things, I think that was it.

Yost: How large was the R&D group within the center?

Bell: Well, that was one of our problems. The director — what was his name? Mel Klein — had been told on his review for the year that . . . . One of the line requirements . . . . You could mark . . . either you listed a set of things that you cared about. There are a certain set of them if you didn't max them out, you're toast. Not toast, but you didn't get all of the perks you could get. One of the ones they put on him was size of the organization. And so he turned to all of his people and said, “Size of the organization.”

When I arrived in 1983 I was number 17 in C3, the research group. By the time I left, I was in the lowest third. So we had tripled in two years, which meant that the people who were told to take care of this new person had been there a year, year and a half. So we had people but we had a difficult time training them and getting them up to speed.

Yost: What led your decision to leave NSA and join TIS?

Bell: Well, when I came down, I had said I either had to be having an awful lot of fun or something needs to happen to keep me going. And by the end I wasn’t having fun, and so . . . . I was just getting the normal increases and I just needed to get out from under it. There were pressures to use part of our budget that weren’t security and I was trying to do budget. And I built the — what was it? — the 1985 Computer Security Program, trying to get all the services and everybody to work together, and all of that. I met a lot of great people from other places but part of our management kind of countermanded some of that stuff and made life hard to live. So that was just no fun.

Yost: I understand Steve Walker started TIS in his garage but then built a building for the company on his property. Is that correct?

Bell: He lived in Glenwood and he had . . . . You know, when he had worked at NSA when he was finishing his degree. And then he worked at NSA — let's see if I've got the right order — he worked at NSA, then he went to DARPA, then he went down to the Pentagon. And when he had done those things, he started in his garage, but he and his then-wife had wanted to move. They had sort of outgrown the house. And they had found this big old house with lots of property and lots of outbuildings, going back to just after or just before the Civil War. And so they bought that and he was going to put it [Ed. his company, Trusted Information Systems] in an old abandoned gas station. Have you been out to see Steve?

Yost: No. I will in November.

Bell: Okay. So when I went to work they were in the rehabbed two-story gas station. And I was the fifth employee. We rapidly started gathering people and it became clear that the building wasn't big enough so Steve went into "let's build a bigger building." We were growing faster than that and we started having . . . . You know Ted Lee was TIS North in Minneapolis, and then we started expanding into — I don't remember the timing — but Los Angeles. We had Los Angeles; and we had my friend Tom Parenty. Oh, you need to go interview him (laughs). It'd be fun for you to interview my friend Tom because he lives in Hong Kong. He was in Berkeley, in San Francisco. But there in Maryland, we were increasing so fast that we were outgrowing the facility, so as they were building the building that initially was just a long building two stories high. After I left they got a lot bigger. So we got some temporary housing, trailer-y things. People didn’t particularly like going into those little trailer-y things. So it was set here in north/south; and we had the trailer; and people back there started naming the computers Toto because they were being sent to Kansas. We outgrew that one. So we got another trailer and put it . . . that was Nebraska. Then we got the big building.

Yost: When you arrived can you describe what the general business model was for the company? What areas, what businesses did it target?

Bell: We were a consulting company and mostly for straight governmental people who were doing government contracts. Steve and Heidi Heiden, Heidi who was a retired colonel, I think, or lieutenant colonel, I don’t remember, in the Army. All of their contacts or most of their contacts were DoD related. So most of our contracts were there, but Steve also had had, from his position in the Pentagon, had a lot of people that were from the commercial companies that sold to the government. In fact, well, he’ll tell you the story . . . . But he always says that he decided to start the company and run this little consulting thing until he could decide which job to take. But he made money the first month. So he never quit. He was wondering, you know, IBM, or go someplace else. So, we ended up . . . .   What we always said was we were the security department that nobody could hire. So we did it for NSA, we did it for different parts of, you know, for DCEC. People would call. One of the things we used to do somewhere in the middle — not when I first got there because we didn’t have enough people — but over time, expanded it and if Motorola was interested in what are the security characteristics we should keep or get rid of in the 68020 or 68030 chip and we’d say “Okay we’ll check —

one month thing and it'll be so much. If you want us to be a part of what we'll give you is . . . . This is what we can help you with." And then we'd have longer term things. So we went from consultancy to having longer-term things. We started winning bigger projects like [Ed. part of] the DARPA Strategic Computing Initiative at the . . . . You know, analyze all the different operating systems that were under consideration. And in particular, that's where we started doing Trusted Mach. So it was mostly, early on, it was like I said, consulting — various people, some industry but an awful lot of sort of a military flavor.

One of the things that we were building on was Trusted Xenix that IBM had come up with. We were one of their biggest customers so when they were ready to terminate the program, terminate the product, they called us up and said "We're going to terminate it." Steve said "Don't make that announcement" and went over and managed to, in 42 days, to arrange a transfer of all of their intellectual property and ownership in Trusted Xenix to us. They called it "Secure Xenix" and Steve renamed it to "Trusted Xenix." That was only for dealing with the paperwork. I was the Corporate Secretary; after being there a year I became Corporate Secretary. So at that point, we started having products. And after I left they started doing the firewalls, which was a big product endeavor.

Yost: So you had a managerial role, but you also had a technical role on Trusted Xenix?

Bell: At Trusted Information Systems we had . . . . Steve was the president and we had four vice presidents. All of us had oversight responsibility as well as technical direction. Nobody sat around and just managed.

Yost: You had quite a dream team of early computer security pioneers there. Yourself, and Marv Schaefer, Steve Lipner, and other notables?

Bell: Lipner came after I was there, but Martha Branstead was there, and Steve Crocker was there. And Steve said at one point — you know we had feelers from a variety of people — and Steve said that he always thought that we could never have all of those people, but he started to reconsider; you know, why not? Could hire all of them.

Yost: Was Trusted Xenix the first certified B2?

Bell: No. Multics was the first. Well, actually, I ended up being the lead to finish that [Ed. the Trusted Xenix evaluation]. You know IBM was so far, and part of what we had to do was to understand the status of the evaluation and what we were facing in what we had to do to finish it up. Part of what we needed . . . we didn't have good model interpretation so I ended up having to write one for that, as well. "Trusted Xenix Interpretation." I wrote it and presented at the NBS Conference. It turned out that UNIX was not enough like Multics that we could just use the same rules. We had to vary it a little bit.

Yost: Did you do much . . . or were you still there when the initial firewall work began?

Bell: No, I had left before that. I left in 1991 and the firewall stuff, I think it was 1990 when they had . . . . So you could interview AT&T's guy, whose name I'm going to have trouble with. The firewall guy. I'll remember later. [Ed. Bill Cheswick and Steve Bellovin.] They had a little talk about the half decade of the firewall in 1990. I was interested in helping staff members here, there, and yonder to understand the basics so that the company could raise the level so that what we at TIS needed to do would be at a slightly higher level. And that just wasn't part of the scheme, doing education things. Now, they changed their mind later. But what I wanted to be able to do was that — teaching at the NSA schoolhouse and going off. So that's when I went off to form my own company for a while. I really wanted a couple, three projects in hand before I left, but that's hard to schedule. When you're working at one company you can, but that's sort of ongoing, [not like] when you're getting ready to start. So I ended up with some training but I also when I told some of my friends back at NSA that I was getting ready to go off on my own, they got all excited and took me over to the people who were doing Dockmaster II, and they ended up writing me a contract to support them to help replace Dockmaster, which was a Multics system with an upgrade. So then I got that contract and . . . that contract followed me. Started at my company, BBND. And then when I went to Galaxy, the contract followed me. When I left Galaxy and went to MITRE the contract followed me. So I worked on Dockmaster II a long time.

Yost: In that roughly half decade at TIS, were there any companies that looked like TIS? Who did TIS compete against? Was it just competing against things being done internally within government versus contracting with TIS or were there other contractors like TIS?

Bell: There were other companies. We, of course, thought we were the best but SCC [Ed. Secure Computing Corporation] did things, Sparta did some things, Arca. I'm not sure I can remember them all, but yes, there were other companies doing some other things out there. And in fact, when people left TIS, they went to some of those other companies. There's a company now called Tresys, T-R-E-S-Y-S, that's run by Frank Mayer, who used to be at TIS. I think Steve Walker helped him get started. And, you know, IBM started doing consulting in this area, and there was a consulting group at DEC. And Paul Karger, after he left DEC, went to IBM where they were doing security work. NRL in the government was doing security work. SDC was doing security work. Don Good down at the University of Texas Austin was doing verification work. After the Computer Security Center, well, the Computer Security Initiative, in general, but the Center and some of those advances, all sorts of people started doing various parts of security around.

Yost: You mention Dockmaster at BBND. Can you discuss that a bit, as well as other work you did for that enterprise from 1991 to 1994?

Bell: Dockmaster was set up as an example system, in the first place, when the center was set up. It was a Multics system, and they had set it up before I had . . . . I mean, they had set it up and part of what they did, and they made use of it in product evaluation because a lot of the information that was traded between evaluators and the companies was company proprietary. So they made use of the security properties to keep that isolated from garden-variety users. What they didn’t do was give everybody a fake security level and category. So you could have experience with using a multilevel system, which would’ve had benefits. People would have understood what you had to go through to do things. Feedback could’ve happened. People could’ve had good ideas about things you needed.

One of the things we found in Trusted Xenix was you needed something provided by the vendor that allowed you to downgrade when you wanted to. All of our systems had been built never to let you do that and so people had found out, found ways to do that that were outside that.

In fact, that system [Ed. Dockmaster] was quite old and Honeywell was getting ready to terminate it so NSA decided that they needed to replace the B2 system to be the new Dockmaster, so they put out a procurement for it. So it was required to be B2, and I was there from the source selection time. I don’t remember if it was much before that, but at least then. The products that were available to try to build this were few and the company that was chosen was BDM. It has now been bought by somebody, I think. I think they got bought by somebody. I don’t remember who. [Ed. TRW.] And their proposal was to use a Data General System that had . . . where the evaluation hadn’t been finished yet. I don’t think they ever did finish it. So we were then faced with a procurement. You know, we had a whole set of features we wanted. We need this product, and it needs to be B2, and also had finished its evaluation. So there was an awful lot of back and forth with trying to figure out how to do this. I think we even had a trip down to Atlanta, or something, where Data General was doing stuff.

Somewhere in the middle, people started thinking to themselves that maybe what we needed to do was to make it web based. Everything's web based. The interface to Multics was all textual — it's type, and things print on the screen. This caused some consternation because people weren't sure exactly how to do it. They decided to move forward cautiously with this and it raised a number of technical issues such as . . . . If you think you understand telnet and its persistent connection, what's this http business? Non-persistent. (Laughs.) What are we going to do about IP — what's it called? IP highjacking, something like that. These are all sorts of issues. Well, every time somebody would raise something like this, they didn't say "What are the security implications?" What kept happening is, BDM kept saying "This is what we're going to do." The government would say "Wait, wait, wait, wait! How we going to do B2 and do this?" I usually ended up being tasked to go figure out what the good news/bad news is with this. So I did all sorts of those kinds of studies. One of the things that was of great concern was whether they were going to, . . . *how* they were going to do modeling and modeling interpretation. They hadn't finished. We had no insight into the modeling that they were doing. In addition to which they were making use of a database — don't even remember. It probably was Oracle, but I'm not positive — and there were no good, wonderful security models that helped you figure out what to do about databases. So I went to the government and said, "This is a risk area." And we talked about it and what we concluded was that without that in the contract what the government needed to do was to have a backup plan, in case they weren't coming through. What I contributed was that since their system was UNIX-based and it was POSIX compliant . . . . So we know what POSIX is. And since their system did SQL, which is sort of standardized — there are varieties, but at least it's sort of standardized — if you pulled the interpretation part out you can say "Here's the model. Here's a generic version of UNIX," and then you could say "Here's a specific UNIX and that's how it fits the generic one." And the same thing for the SQL. So what we ended up deciding to have me do was to have me quietly, without telling anybody, build a generic modeling interpretation that would work for anything that claims to be POSIX so that if they didn't come through, we had the fall-back position of saying "Okay, so we've got this far, can't we hook it here?" You know, we've got modeling, we've got interpretation here until we get this other piece. So at some point — don't remember why we decided we weren't going to keep it secret anymore. But I submitted it, and had a paper that talked about how everybody is doing POSIX, you know, is most of the way home on an interpretation. It somewhat built on Trusted Xenix because it's UNIX, too. I'm not sure if it was POSIX compliant but it came fairly close. [Ed. The paper was "Generic Modeling Interpretation."]

Yost: You were the chief scientist at Galaxy Computer Services for a short time.

Bell: Yes. And mostly what I did there . . . . Well, let me see. You know, I forgot something. When I left TIS, I ended up getting, shortly after that, not just education [Ed. tasks]. The other thing that I got . . . . Blaine Burnham, at the time, was the guy in charge of the Rainbow Series. And he said “What I need is I want you to go quietly into a room and write me an updated abridged book [Ed. to update the TCSEC] that . . . .” He said, “I don’t want anybody else to know about it. I don’t want anybody else helping. We just need to write something that’s good.” Outside events changed that from quietly where we’re in a corner, to a group of five, to a group of seventeen, to fights with the European criteria, and the North American criteria, and it exploded out of all of it. So one of the things I was doing besides Dockmaster was working on the follow-up, whose name kept changing. So while I was at one of the . . . . The funding method for that work on the criteria was that they gave a contract to Galaxy, subbed to me. So I was working with them on working on the update of The Orange Book, and traveling to Albuquerque a lot to work in Dave’s house [Ed. Dave Bailey]. I was living here and Kate and I were taking a walk. We were walking along and talking about the nice fall colors, and without thinking about it I said, “You know, I want a job.” Because when you work consulting for yourself you have to put in 40 hours to get paid 40 hours. Then you have to put in 20 or 25 to run the business. I’d done it for two years and I was worn down. So I talked to a number of people and ended up going to Galaxy. What I did was continue with the projects I was on. I went from a subcontract to an employee. There were a couple of other things that I did, but those things, those same two projects.

Yost: And Mitretek?

Bell: It was . . . . You know, my wife works at MITRE. I met her in 1992 and got married in 1993. She kept encouraging me to come back to MITRE. I had worked at

MITRE Bedford that was having a fight with MITRE McLean for who was doing the best work, and so on, and MITRE Bedford lost. MITRE McLean is sort of in charge. It turns out now that most of MITRE is here at McLean, not at Bedford or any place else. She kept urging me to go back to MITRE and in 1995 I applied to MITRE in the summertime. About a couple of weeks after I got there they announced, MITRE announced, that it was being split in two. MITRE's Air Force sponsors . . . MITRE's sponsors when I was there the first eleven years, they went back and forth. They would say "Do all DoD." "Oh, cover yourself a little bit with some other way because we don't have enough money." "No, do all DoD." "Cover yourself." So it went back and forth between diversifying and concentrating of the core job. Well, we had gotten one of these concentrate-on-your-core jobs, and the management looked at it and said . . . . They ended up at a summit, and part of management said and "We'll do this other stuff," and the Air Force or somebody said "No, we don't want you to do that." So what MITRE management decided to do (in Marv Schaefer's phrase) "infinitesimal wisdom" was that the work that they were thinking that they wanted to do commercially, what they were going to do was split the company into MITRE and what turned into Mitretek Systems. And then part of Mitretek Systems grew to be another nonprofit, but just not an Government-sponsored FFRDC for the Air Force, and FAA, and some other things. Then they split off a for-profit company to go work for for-profit people [Ed. Concept 5].

So I arrived at MITRE and found out that in six months we were going to be split. Kate was going to stay at MITRE; *I* was going to be at Mitretek. Well, there were 25 and 30 married couples that we learned about that had that same hands-across-the-sea. So for six months I worked for MITRE, and then I worked for Mitretek Systems. I was the chief technical guy in a MITRE department that did security, J24 I think it was, and I still had my Dockmaster work going. But we also acted as kind of an internal consulting organization for various government folks. So one time, the FCC called us up and said "Come and see about our security." They had great security. I went off to see the — what's it called? — the Overseas Investment Center; no, that's not quite right. [Ed. Overseas Private Investment Corporation] They had a real problem. They used to send people out with confidential stuff on laptops with not even passwords, going to foreign countries and going through customs. So, different people.

So somebody else did Unemployment Insurance. They were doing a pilot project where they needed to exchange data between IBM mainframes and PC kinds of things. They needed to ship things out that needed to be encrypted while it was going through, and it needed to be triple DES. Do you know how expensive triple DES is on an IBM?

Yost: No.

Bell: $30,000. And, you know, they took the "these are what you need to do" and we looked at it and we said "This is what you need to do" and they said "Don't we need to be public keys?" "You only have 12 places." (Laughs.) It was easier just to get it out there. You don't want that overhead. So they said "Find us sources." We went out and found sources. $30,000. "Try again." And that was it. So, gosh, I can't remember them all but part of that department was continuing to do evaluations supporting of NSA.

I wasn’t involved in that. I tended to be involved in the other things that tended to come to us. I also ran MITRE Independent Research on the inside. Everybody — remember, this is 1995-2000 — everybody that went off to do public key stuff, they’d go off to do it and they’d get stuck on the early stuff, they’d never get to something substantive. We made a proposal to get to the substantive stuff and we didn’t get to it either, but for a different reason. We got roadblocks on what we were trying to do. What we wanted to do was to do an example that worked into real business, like time cards, or expense vouchers, and the hard part was how you going to hook that into the actual back office stuff. We made good progress until I left. I mean, we were making progress. We weren’t going to reach all of our goals, but we were making okay progress when I left there and went off to EDS.

Yost: And you were Chief Scientist, Information Assurance?

Bell: Yes. They had a Center of Excellence. There was an Air Force general who left to be a senior vice president for the government area over there near the airport. Al Edmonds? Sounds right. And what I was told was that when he came there and met the people he said “Where’s my information assurance center of excellence?” And they looked around and said “We don’t have one of those.” “Might want to get one.” So one of the guys that I had vetted, a retired Army colonel, Dain Gary, who had been at NSA . . . . I was polling around and he said that they’ve got this organization. They need somebody senior. So I went over and had one interview. Came back a second time and they made me a verbal offer, which . . . . I was so stunned about I didn’t ask them to write it down. Guy that hired me, Rick Rosenberg, came into his office and he stood up. He was about 6’5”, probably 280. He was a former NFL football player whose career ended his first year with a career-ending injury. Now he was a senior vice president here at EDS. He was really big. So we sat and talked a little bit. I came back the second time and he made me the verbal offer. So I went over there and my job was to work on individual projects, as well as help nurture the other people there in the organization to get better at their security. Most of the people had . . . . You know, the managers were at a mid-level. Most of the others were fairly junior.

Yost: Was EDS doing much high-level security and government contracting or was this more in the commercial sector?

Bell: This particular part of EDS was part of their federal services branch, whatever it was called. So we did work with INS. What actually happened while I was there was I got there in 1999 — think it was 1999 — and suddenly there came a call for EDS to send a bunch of people down to Department of Energy. The CIO there was John Gilligan and he arrived just in time to get a failing mark for Y2K. After he scrambled getting ready for Y2K, he decided he needed to work on security. So he called in a number of people with companies he already had on contract and tried to figure out what was to be done about getting some good security policies for the Department of Energy. Department of Energy is not homogeneous. It was jammed together under Carter and it’s still not one agency.

So he said, “I know you guys are all used to dealing with contracts, but you need to work with me here and try and do this stuff.” So I ended up regularly coming downtown to the Forrestal Building every single day, working there and he was trying to get two deputies for covering different things. One of them was a technical advisor on security matters, and he had it posted, but even if somebody showed up they weren’t going to get it [Ed filled quickly]. One of our salesmen, former SEAL, came to EDS guys and said “We should propose David to be that guy.” So we talked about it and decided to do that. And so he [Ed. Gilligan] said maybe I should have him come in, he wants to talk. So I went . . . down at my office buried in the bowels of the basement, went to the top floor and went in to talk to him. Well, it turns out he was a grad student at Case Western when I was there doing the modeling, and he was there when I told the people they had done their math wrong. So it was old home week because he’d gone to SDC and he knew a number of people there and we talked about things we knew. So I ended up his technical advisor and an EDS employee at the same time while we were trying to figure out how to build a policy that would apply all across the various things.

One of the things, one of the objections to this building a policy, people out in the field said “You guys back in Washington don’t understand what we do.” And he concluded that they were right, so we called it the “CIO World Tour.” We went off for two weeks, came back for a couple of days, and went back for another two weeks of visiting various sites talking to them about what they were doing and how they were doing it, gathering information, talking to them about what we were doing. Very illuminating. Most of the time, most of the people were hostile. “People from headquarters don’t understand us and we’re smarter than you anyway.” But not all. And we learned all sorts of things. We got tours of underground. You know, “That’s where the ring is, right out there,” you know, atomic particles. And we worked hard on getting that put together. I was working hard doing all sorts of things, then my son decided he was going to drop out of college and I tried to get him not to, but he was going to get a job at RealNetworks. This was in Seattle. He decided to do . . . . One of the things I said was, “You know that loan that I’m paying on? I’m not doing that anymore. You know that money that I send you for school? I’m not sending that anymore. That money I’m sending for you to live on? I’m not sending that anymore.” Anyway, it worked out great. He’s now at Microsoft. He’s a lead interface designer for Microsoft phone. But, I concluded with that much money saved, then maybe I could retire. So I did the calculations and I did. So I was working supporting Gilligan, and then I retired in 2000.

Yost: And then you came back from retirement, as a consultant for the Air Force?

Bell: Yes. What happened was that I had calculated if I could retire, based on money and things like that. And I assumed that the stock market and the interest rates would all be the same. After the attacks on September 11, I decided that some of my assumptions were no longer true. I couldn’t be certain things were going to stay the same. I decided that I needed to unretire so that somebody else could go drive a truck or shoot a rifle. I polled around and found out that by that time Gillian had gone off to the Air Force, shortly to be named Air Force CIO. I also talked to MITRE, and also talked to somebody else — I forgot who — and ended up going to support the Air Force. I had figured that he would want me to become a government employee. We talked over the phone and he told me that to hire me as a government employee would take over six months but that if I subcontracted to the people he had a couple contracts with, I could come to work tomorrow. It was clear that he preferred option two so I went and did that for about a year.

Yost: Then back to MITRE?

Bell: Then I went back to MITRE.

Yost: Was that doing computer security work at MITRE between 2002 and 2005?

Bell: Yes. I went back into the group at MITRE that does security. It's the security division.

Yost: When I interviewed Roger Schell, he offered a pretty bleak picture of where he thought security was today. With you coming out of the same kind of high assurance work, do you hold a similar view or a different view?

Bell: I think it's pretty bleak. The paper I did in 2005, right after I retired, what I said then was that we have a few really good resources but people won't use them because they've been blocked out procedurally, policy, a variety of things of that sort that keeps them from being allowed. There are some people who want to use them but they don't believe they'll get it through their certification and accreditation authority, so they won't try. And what we are doing is putting systems, Windows systems, plain old UNIX systems, even Trusted Solaris systems are what The Orange Book would call C2 or B1 systems.   Which are perfectly adequate to keep cooperating colleagues apart. We are putting those between networks of totally different security levels or confidence levels. So we are putting our weakest systems at our most critical junctions. Where we have the option, now, because of the way we network . . . . It might be better for every computer, or most computers, to be extremely high assurance. But one needs a transition and you can't snap your fingers and have it be that way. But because we are in networks with isolation points between the website the company offers and the back office, between unclassified DoD network and a classified network, one likes to think of those as classified and separate. In real life, there are these limited connections in between. And those limited connections in between are, by nature, multilevel. Front office, public, something is sitting in the middle trying to keep things from going in the wrong direction there. Because that's the case, it would be fairly simple to put the functionality you need to stop whatever you want to stop, but do it with high assurance if people were willing to go ahead and do it.

I know less about the commercial world, but the commercial world is not buying the products that are available and have been available for 10 or 20 years. The government purchasers are in a position where the approvers won't approve it, so the program managers and the contractors won't propose it. It's not a good situation. In my follow-up paper, "Looking Back: Addendum," I pointed out that a large number of government initiatives had the effect of stymieing or discouraging people from making use of good security with the promise that they'll provide you something even better in the future. They promised it in 2005 and they cancelled almost every one of those programs. I don't know what the current promise is because I'm retired and I'm not keeping up but as Marv Schaefer used to say, "I fear for the republic."

Yost: Finally, are there any questions I haven't asked or topics I haven't covered that you think are important to cover before we conclude?

Bell: Well, I guess we've covered it a bit but to deny that you have a problem doesn't make it go away. There are people who say we are perfectly safe and if they don't know, we shouldn't listen to them. If they do know, they're lying to us. We have real problems and as Carl Landwehr once said, for 20 years, 30 years, we've known what to do and we actually have some products that can help us but we're not doing it.

Yost: Okay, well, thank you. This has been very helpful.

Bell: Good. Glad to do it.