

An Interview with  
CARL E. LANDWEHR

OH 436

Conducted by Jeffrey R. Yost

on

21 April 2014

Computer Security History Project

McLean, Virginia

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

## Carl E. Landwehr Interview

21 April 2014

Oral History 436

### Abstract

Computer security pioneer Carl Landwehr discusses his educational training (Ph.D. University of Michigan), his research as computer scientist/supervisory computer scientist at the Naval Research Laboratory in the second half of the 1970s, 1980s and 1990s, and subsequent work as a research program officer for computer security at the National Science Foundation (over two separate tenures) and IARPA (where he served as a Division Chief). Among the topics discussed are the Secure Military Message System Project, survey work analyzing early security models, his work on application-based security models, and the role of federal research programs in advancing the field of computer security.

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 at the University of Minnesota, and I'm here to day on April 21, 2014 with Carl Landwehr in McLean, Virginia. This is for CBI's NSF-sponsored project, Building an Infrastructure for Computer Security History. Carl, can you begin by just giving me a little biographical information, when and where you were born?

Landwehr: I was born in Evanston, Illinois, on September 3, 1946. I spent the first six or seven years of my life in Northbrook, Illinois, where my father had grown up. We moved to Elmhurst in 1953 and I continued to live there until I went away to college and started a career.

Yost: Can you describe yourself as a student, your pre-college days? What were your interests and what people in your life were influential to you?

Landwehr: Pre-college?

Yost: Yes.

Landwehr: I had a number of elementary school teachers and other teachers that I remember fondly, who I think were very good influences, most of them in Elmhurst. At junior high school, there was an English teacher who was very good, named Miss Bingham, at the time. Then in high school, I had a good series of honors English classes, and mathematics classes. York High School provided a very good education at the time, I

thought. Freshman math teacher, Mr. Zwoyer, I believe was his name, was very good; and there were others. But actually, I was interested in a lot of activities. I edited the school paper in my senior year there; Ms. Eleanor Davis was the journalism teacher.

Yost: Started as an editor early on.

Landwehr: I did. My father was actually an editor of weekly newspapers on the west side of Chicago, and ended his career editing the *Southtown Economist*, which became a daily when the *Chicago Daily News* folded. And my mother was an English teacher, actually, as well so I have a lot of that in my background. My older brother, who's a statistician — worked at Bell Labs for many years — is actually probably the mathematician of the family. And I have a younger sister who's a lawyer.

Yost: How did you decide to go to Yale University and when you first went there, had you already set on studying engineering and applied mathematics or did that come later?

Landwehr: I went there because my brother was already there and I had a couple of cousins there. I actually interviewed a lot of places, some other places, and might've gone elsewhere but in the end, I think, my father considered it simpler for us to be in the same place on the same schedule, and so that's how I wound up there. And in our family I had a number of cousins, all in the Chicago area, but the oldest of those actually was the first one to have gone to Yale. His father was in the nickel business and he studied metallurgy, it was the late 1950s, early 1960s, and there was a lot of emphasis on science and

engineering. When I was in high school, actually I went to a summer institute at Northwestern, National High School Summer Institute, something like that. It was known as the Cherub Program and that was funded, I think, probably by NSF or somebody at the time to try to encourage people to go into engineering. So I think I started out in the engineering program but engineering at Yale — Yale has had a mixed history with science, in a way, and at the time, engineering was no longer a school it was a department. It was Engineering and Applied Science, and so the curriculum was pretty science oriented. But they did have a professor who was very influential, I'd say, on several of us there, named Bob Rosin, who was a graduate of MIT and Michigan, actually, the Communication Sciences program at Michigan. And he taught undergraduate computing courses. It wasn't, I guess, the first programming course I had but maybe the second, and I basically took as many courses as I could find that were computing related. But there was no computer science department at that time at Yale, They made an effort to recruit some computer scientists, but I think they didn't offer enough; they felt that it was sufficient that they were coming to Yale, they didn't offer a lot else. So no department got created until later when they finally recruited Alan Perlis and started a computer science program. But Bob Rosin was a strong influence and we did a lot of work at the computing center there.

Yost: Do you recall what systems you used there?

Landwehr: Yes, I actually have a core plane from an old [IBM] 709. They were replacing a 709 with a [IBM] 7094, and they had a, was it [IBM] 7040, [IBM] 7090 direct

coupled system, which had batch programming and these giant disks you could watch move, and it was a lot of fun there. And I had friends in the engineering program. One I went to high school with, Dean Kloker, also was with that program; he is now in Minneapolis. So we enjoyed that a lot and I think partly because of Bob Rosin's influence, I was interested in the program at Michigan, which at that time was not actually a computer science program. It was called communication sciences. Michigan had a very fragmented computing situation at the time. There was computing at a number of different places. In the electrical engineering department there was some; in the industrial engineering department was where most of the graphics was going on; there was the computing center, which developed the Michigan Time-sharing System, and that interested me. Anyway, the program itself was an interdisciplinary program, which included electrical engineering aspects. It was about information processing, really, in all kinds of systems, so it included information processing in human systems, so biological systems. We had some introductory courses in psychology, and in linguistics, and philosophy. It was a fascinating program to be part of but it also was described by some as a mile wide and an inch deep; you had to specialize when you were going to do the dissertation research after you did this. This was a program that was the child of; I think the creation, really, of John Holland. And Art Burks was there at the time, also.

Yost: Bernie Galler?

Landwehr: Yes, Bernie Galler.

Yost: Bernie was a good friend of our institute and a personal friend. I worked with him on a software history project funded by NSF.

Landwehr: Yes indeed, Bernie was very interested in history. His passing was a blow. Bernie was a very popular professor, and he and Bruce Arden were actually chairs of my dissertation committee. But in fact, my dissertation really was about numerical simulations of queuing theory and so I worked actually with Ralph Disney, who was in the industrial engineering department; queuing guy. So what I ended up doing in grad school was working for the MERIT Computer Network. I think Michigan wasn't willing to pony up the money for an ARPANET IMP, and so they were going to build their own, and they did. That gave me an opportunity to participate in developing an operating system for a packet processing machine, set up with a number of people. I worked at the computing center doing that and my dissertation was partly motivated by that.

Yost: Going back for a moment, you also worked at Bell Labs before grad school, I believe in the summer of 1967 . . .

Landwehr: Yes, that's right.

Yost: . . . as a programmer? Can you describe that experience?

Landwehr: That was interesting. That was a summer job and it was the first summer job I had that wasn't a manual labor-type summer job. I had worked at Jewel Food Company

Warehouses before that. But that job — I think I got that through the Yale summer job application process or something — but it was the summer that Bell Labs had just opened up their Indian Hills facility in Naperville. Brand new, and at that time, it was about half empty because I think they had trouble getting people to move out from New Jersey. But they had an IBM 360 model 67, which was trying to run TSS and I was just, you know, I was very junior, of course, and I was working on some programs involving personnel data. But there were sometimes staff meetings, which there was no need for me to attend. So I remember one afternoon I actually had the entire computing hardware to myself and was really in charge of booting this thing up from scratch, which was somewhat daunting but I did it all right. But, yes, TSS was being debugged at that point, and I think the operational stuff was OS/360; that's where I learned whatever I did learn about JCL, a horrible command language that they created for that. People I worked with included Russ Archer, who was head of the group; I believe Nick Martellato was his boss. Other names escape me.

Yost: Can you expand a little bit on what specifically you worked on with MERIT?

Landwehr: So MERIT produced a communications device which was designed to sit in front of each host system. There were three host systems: Wayne State, University of Michigan, and Michigan State. They all had, let's see; Michigan State, I think had Control Data equipment. Michigan had IBM equipment. Michigan State I think also had IBM equipment but Michigan was running MTS and I think Wayne was running — I'm not 100 percent sure what Wayne was running. Anyway, the idea was that this device;



there was a packet switching device called a Communications Computer (CC) that sat in front of each one and then there was some—

Yost: So the equivalent of the IMPs with ARPANET?

Landwehr: Yes, the equivalent of that. And so there was development of the operating system for the CC; I don't know if I generated code for that, but we had designed and discussed it. In fact we decided to use semaphores. Someone else decided that [Brian Read and Al Cocanower, probably], it wasn't my decision. So then we built everything around that sort of coordination structure. But the part I worked on was primarily; there was a piece of software, a device support routine, which in MTS talked to that device. And so I wrote the code for that; both the main code on that host [MTS] side and also on the MERIT Communications Computer [CC] side. I had to write the code on both sides of that interface.

Yost: I don't suppose there was discussion of computer security issues with that network?

Landwehr: Security was actually an issue in the air around there, and at that time, of course, to debug the operating system and so on, I mean, there was one computing system for the campus. And so virtual machines were used, in fact that was really — from my perspective at the time anyway — why they were created. It was so that you could debug the operating system without taking the system over and running it by yourself. Michigan

had a very well-developed accounting system for rationing time and students had rations of time they could use on courses, and so occasionally students would try to get extra resources one way or another, and sometimes people would play games. So there were definitely thoughts about security, and in the context of MERIT [pause]

Yost: So that was with the Michigan Terminal System?

Landwehr: Yes, that was with MTS. In the context of MERIT, I don't remember explicit security discussions,

Yost: And what year did you start in the Michigan computer center?

Landwehr: I started Michigan in the fall of 1968 and left in the fall of 1974. So, I didn't start working for MERIT until, I think; let's see, the first summer I spent actually at Lawrence Livermore Lab, the summer of 1969, and then the summer of 1970 I think I started working for MERIT. It was either work for MERIT or to take a job with the tennis coach teaching tennis in Ohio someplace.

Yost: And in 1969, for Livermore, was that a programming position?

Landwehr: Like a large number of people at that time, I was confronted with the possibility of being drafted, and I decided if I didn't get drafted I was going to spend the summer in California. So I applied for jobs at the RAND Corporation, I think, and

Livermore, and other places like that. I was able to, in the end, avoid the draft because of my eyes, actually. And once I got that deferment I said that's it. So I spent the summer there working with Control Data equipment at Livermore. They were, I think, just getting a CDC 7600 in and they had a lot of CDC 6600s. What was interesting there was it seemed very backwards, compared to Michigan. They had these amazing online card punches that would suck in an entire box of cards in a matter of seconds. But they needed them because they didn't have enough storage to store their files overnight, so at the end of the day, they would punch out their cards and then in the morning they would read them in. Actually, that's what I ended up doing there: support for the online card punch.

Yost: In applying to Rand, were you aware of the work of Willis Ware and his early work in computer security?

Landwehr: No, not at that time.

Yost: In 1967, he and NSA's Bernard Peters wrote a paper on multilevel security with the proliferation of time-sharing.

Landwehr: No, I wasn't really aware of that. No, that was strictly; actually, one of my roommates, Rich Jagacinski, did get a job at RAND and we used to meet on the weekends in the National Parks.

Yost: You completed your dissertation the summer of 1974, and it was entitled “Load Sharing in Computer Networks: A Queuing Model.” Can you describe that dissertation?

Landwehr: I suspect that you can count the number of people who’ve actually read that dissertation on the fingers of two hands, at most, [Laughs.] It built on work of other Michigan grad students who had developed queuing models. Kip Moore had developed queuing models primarily for optimizing paging drums, as I remember, on the system and other people built on that. Vic Wallace had developed some numerical analysis software. The issue at that time — this is before the famous queuing papers from Baskett, Chandy, Muntz and Palacios about how to compose queues, if you put a lot of constraints on them. The proclaimed reason for these computing efforts, including the ARPANET was to share these expensive computers that weren’t located in too many places, so the idea was you’re going to do load sharing. So I was trying to do that modeling, and also at that time, time-sharing was beginning to take over from batch, but the load sharing really meant, you know, people thought well, I’ll send a batch job over there and have it done. So I created a model where there was both a time-sharing queuing component and a batch sharing queuing component for a network of only three systems, which is what MERIT had, and then try to model the queuing behavior of that. I did publish one paper out of that, eventually, with Erol Gelenbe who was on the committee as well, was still around and working. So it was for me, a challenging thing to do, and I was happy to find another area to explore afterwards, I guess is the right way to put it. In working on MERIT, it was interesting because I learned a lot about how networking was going to work.

Yost: Can you talk a little bit about your mentors on that project?

Landwehr: As I had mentioned already, I think Ralph Disney was probably the primary one from the standpoint of the queuing theory and the numerical analysis of queues. And Bernie Galler was a steadfast advisor; and Bruce Arden helped out, but Bruce at that time took a sabbatical in France for a year, maybe even longer, at Grenoble and so he was absent for a fair chunk of the work. Al Cocanower, who was full time on the MERIT project, also served on the committee and advised me. So a lot of it I can remember trying to work out various stochastic formulas for that stuff and, you know, the idea of being able to measure the system and measure the performance seemed to me, and still seems to me, that that's an important thing to be able to do. Other people have been more successful at queuing theory than I was.

Yost: When you entered graduate school, what were you thinking in terms of career, and were your thoughts the same as you were finishing graduate school?

Landwehr: I don't think I really thought very much about it; I thought well, I enjoyed working with computers and I was happy to do that. Over time in grad school you realize that you can't remain a grad student forever, that's not the idea. In fact, I remember Bob Rosin telling people your job at grad school is to get out. I enjoyed Ann Arbor a lot; I met my wife there, and we were married there; well, we weren't, we were married in Chicago, but we became man and wife while we were there. So it was a happy time, but

those times you have to finish. And when it came time to finish, that's what I was working on so that's what I did.

Yost: So the thought of industry, government, or academia—

Landwehr: I think the advice I got from Bernie was if you think you might want to take an academic job, you should do that first because you'll have a harder time coming back. So the strategy was to look for academic jobs. I did look at a couple on the west coast that were involved with support for the ILLIAC out there, but that would've been a different path. But mostly, I interviewed at different universities, and at the same time my wife was finishing her degree in environmental health sciences; she did water quality modeling and she did not want to do an academic job. So I was trying to keep my eyes open for jobs for her and we wound up at Purdue through that process. Unfortunately, when she got to Purdue, all the offers she actually got were grad student level jobs, so after about nine or 12 months we said well, this didn't work, let's try it again. I went back again and talked with Bernie and others, and he said places for two-career couples that look good are the following kinds of cities, and Washington was one of them. That's where we wound up and that's worked out very well.

Yost: Peter Denning, of course, was there at Purdue as well as, I believe, Dorothy Denning was in her last year as a grad student and became a faculty member the year you were there.

Landwehr: Yes.

Yost: Peter had written in 1972, an important article with Scott Graham entitled “Protection Principles in Practice.” Did either or both of the Dennings have an influence on you in thinking about computer security discipline?

Landwehr: I had an interest in security while I was still at Michigan. Not motivated particularly by MERIT, but just I was interested in it. And then I was not, certainly, on Dorothy’s committee, but I think I did sit in on the oral and heard what was going on there, and I thought that was pretty interesting. So that definitely attracted my attention. One semester at Purdue, I taught the operating systems course that Peter had leant me his notes for, and I taught it as best I could, based on that and so that included some of the security stuff, I believe, although it’s been a long time now.

Yost: In 1976, you published an article on load factors. Can you discuss that article?

Landwehr: Remind me the title. [Laughs.] Is this the one on queuing theory?

Yost: I think it was.

Landwehr: “*An Endogenous*,” probably, whatever it was.

Yost: Yes.

Landwehr: So that was, you know, these days my impression of the graduate education is you go to grad school, write a bunch of papers, and eventually turn it into a dissertation. At that time, the approach was you went to grad school, did some research, finished the dissertation, and then wrote papers. So I basically got one paper out of my dissertation, and I think that's probably it. Erol Gelenbe, I think, was very helpful; that conference was held at Harvard and it actually seemed to get a nice reception at the time but it wasn't holding my interest.

Yost: You mentioned that Purdue wasn't working out for you and your wife.

Landwehr: Yes.

Yost: So you went back on the job market. Can you tell me how you wound up at the Computer Sciences Corporation?

Landwehr: Well, the other experience I had at Purdue at the time, really, was I felt that — I enjoyed the teaching, but the teaching, I found, if you let it, was a consuming activity. Even though I was only there for a year and a half, it was also clear the university was constantly — at least it seemed to me at the time — sending out those newsletters saying so-and-so's gotten this grant and so-and-so's gotten this grant. One of my grad student colleagues, Shi-Bing Yao, was there as well and he was clearly getting into this game and doing well. I thought, "This isn't what I thought it would be," and so,



you know, it was clearly a very entrepreneurial environment, not what I imagined a scholarly environment was, exactly. It was clearly a competition among the faculty members in a way that I hadn't been aware of when I was in grad school, at least. And so I wasn't that enchanted with it, to tell the truth. And I also felt it was far removed from the things—a lot of it got very abstract, away from the fundamentals of computing that had motivated me. So when I interviewed the next time, really, we did it sort of the other way around, and my wife looked around and found there were some good opportunities in Washington, and I said fine, so I'll look for jobs in Washington. And so I interviewed with several places, including MITRE, but it was the part of MITRE that does the air traffic control systems; and a couple of time-sharing companies, I think I might remember their names in a while but it doesn't really matter; and in the end, I had an interview with the part of the Navy that actually built the operational message systems for the Navy. I decided what I really wanted at this point was to go completely as far as I could at that point in my career towards the practical side. And so I took a job, which was really a job with Computer Sciences Corporation, which had a contract to help these guys who built the communication systems, you know, with whatever research they needed to accompany that.

Yost: Computer Sciences Corporation was a large computer services company that really targeted the federal government.

Landwehr: Yes, right. So, the company offices were in Falls Church but I never went there. It was down in the Navy Yard in a rather grubby facility at the time, which had a

Top Secret clearance required to get into it, and the Navy Yard has now been redeveloped a good bit, but at that time, it was a pretty dicey neighborhood. I worked in an office with three other guys, one of whom was a retired Navy chief who had served on submarines in World War II, and the other was a high school dropout — he was a dropout in the sense I think he was probably thrown out. He was obstreperous, but he was actually a very bright programmer and so he had written a lot of the code for these systems. It was called the LDMX initially, and renamed NAVCOMPARS [NAVal COMmunications Processing And Routing System]. Those systems were real distributed message systems that the Navy ran at the time.

Yost: You moved in a short amount of time, from CSC to the Naval Research Laboratory. Can you tell me about how that came about?

Landwehr: Having chosen the path I did, and working for a while, it was clear to me that although I was learning things, it also wasn't exactly the right environment. And then I discovered that one of my colleagues from the MERIT project was working at the Naval Research Laboratory, this was Connie Heitmeyer, and she had been at Michigan, as well. I knew her from there, although we had not been in close touch since then. Actually, I think this connection came about because when I was still working in the Navy Yard, for an organization called NAVCOSSACT, NAVal COMmand System Support ACTivity, there were security issues brought up there. They had Univac 1108 systems and it turns out some people from NRL had been hired to do, I think, penetration testing for that. And there were claims made, they wanted to be able to share information, they wanted to be

able to use the system to do some degree of multi-level secure program operation, and Univac was happy to attest that this was a secure thing to do, but it wasn't so clear that it was. And so there were meetings about security and there was a guy in what was then the Naval electronic systems command named H. O. Lubbes. I remember being in meetings where his name was brought up, although I hadn't met him, I think, at that point, while I was in the Navy Yard. Subsequently, I worked with him at NRL, as he was a sponsor, and then eventually he came over and headed the group. So there were connections that I made in the Navy Yard that were interesting, and I actually wrote some security memos at that time, because I'd been influenced by the experience at Michigan, that I really thought there was a chance for having high assurance secure separation using the kinds of mechanisms that were in time-sharing systems at the time, if you just did it right.

Yost: So it was really that time at CSC that you really solidified a growing focus for you on security —

Landwehr: Well, that got my interest. I mean, I really thought of myself, and I think still, as more of a software engineer. I resisted going into security as a real complete focus for quite a while, even after I was at NRL. When I was at NRL, my first work was with Connie, but we were actually working on simulating communications protocols for satellite communications, which the Navy was beginning to use more of. And ARPA network was really catching hold and so there was a question, could you put an ARPANET IMP on a ship, or something like this. So we started working with that, and with modeling the Aloha protocol, and the demand access protocols for that stuff. So my

first couple years at NRL were really focused on satellite communication simulations that I developed a fairly significant program to do that, which turned out got used a lot longer than I realized, I think, by other people who took it over. I did write a paper about some of that work, because we did it in SIMULA — maybe I'm getting ahead of the story — but in any case, the transition to NRL was a good one, even though the — well, I think from every standpoint. It turned out that the group that I interviewed with, where Connie was, was actually a staff element at that point, to the Director of the Communication Sciences Division which actually seemed poetically correct, since my degree was in computer and communication sciences. There was sort of an *enfant terrible*, by the name of Bruce Wald, who was the head of that division. In the visitor's office, they would have photographs of all the superintendents and they would all have a tie on and shirt, except for Bruce Wald, who had long red hair. And Bruce had hired what turned out to be one of my Yale classmates, John Shore, as the head of this information systems staff. I didn't know John at Yale, but one of my classmates did know him. John was a ball of fire, and he had come down perhaps taking a job there to avoid the draft, I don't know. He had gotten his Ph.D. in physics at, I think, Catholic University and was an aggressive and bright guy, and attracted a good group of people.

Yost: Roughly what was the size of the information systems staff at that time?

Landwehr: Not big. I would guess somewhere between probably 10, something like that, and 20. But we had Dave Parnas as a consultant, and Dave Weiss was one of my colleagues, who's now at Iowa State after a long, distinguished career. Frank Manola was

doing database work; Stan Wilson actually was doing security work and he later became a branch chief, who I worked under, and was a very good guy. And there was Connie, and Honey Elovitz was [pause]

Yost: Had Connie done anything in the computer security field before you got there?

Landwehr: No. At the time, software engineering was really seen as the core, I think, of that group and that's what Dave Parnas was consulting in, and Kathryn Heninger and Paul Clements, other people joined later, and we became a branch at some point, maybe after John left. No, I think we became a branch first and then John left. But anyway, that was, you know, security was somewhat off to the side and think that was one reason I thought this was too narrow a specialization at the time. But I was interested and it turned out that there was work, because, H. O. Lubbes in fact, had signed up to sponsored the development of the SCOMP. He was a contract manager, he needed technical people to oversee it, and that was one of the roles that NRL typically had. And so that got me involved in that, and that provided a stimulating and interesting, separate set of colleagues from outside the lab.

Yost: Can you give a bit of an overview of how work was organized within the lab, and to what level was their time for some more basic research versus applied systems for Naval operations?

Landwehr: NRL is, I think, and probably remains a somewhat unusual lab within the government. I don't remember the overall budget at the time, but in general, it gets a certain amount of money from the Office of Naval Research that's referred to as "six-one (6.1)", which means in the defense line of things, that's basic research funding. And that was the money that people competed for internally. So if you could get 6.1 funding for your project that meant that you could pretty much decide what the project was and you could take it whatever direction it went. You were subject to annual reviews but it wasn't too onerous. I wanted to get that work; and I got a little bit of it I think at one point for some security modeling work. But the other two-thirds of it — 6.1 was maybe one-third of the lab's budget — the other two-thirds of it really came from outside sources, either within the Navy, or some of it came from other agencies from NASA or NIH, because we had a broad range of science going on, as there is today I think, at the lab. So I think one way that the management decided what to put 6.1 money on is, if they thought there wasn't any other way to fund that work, then they would fund it with the 6.1. If they thought you could get outside support, then they didn't. That rankled a bit because in fact, they figured I could get outside support and I did, in general, for all that stuff. But at the same time, there was a prestige associated with the 6.1 work which was not quite there for the sponsored work, and so that was a source of minor irritation, perhaps. But actually, the early work that I did there, I really did have a lot of freedom, and the time was made, and I purposefully took advantage of it, in a way. One of the things that I saw when I looked, for example, at Peter Denning's career, I thought he had done actually very well writing an article for *Computing Surveys* on operating systems and paging algorithms. I thought that was a very valuable paper because a lot of people could read it

and learn from it. That was, I think, one of Peter's real strengths was the ability to communicate in writing. So I thought; I could see this as I was exposed to the development of various computing systems trying to meet security requirements. There were a lot of these efforts going on. There was a lot of gray literature, if you will, that wasn't out there that an academic, where I'd been, wouldn't really be likely to see much of. It wasn't peer reviewed but actually could be highly technical in many ways. So I thought there was a valuable service to be performed in trying to collect that information and make it available in different ways. And I could also see that an academic wasn't going to be rewarded for providing a survey paper of anything; that isn't the way the reward system works. Whereas at NRL, that was okay, that was fine. So I was allowed the time to do things like develop the formal models paper and the best available technologies for security paper, really, they'd written some nice things.

Yost: That formal models paper was really bringing an overview of the field, a look at security and placing it in the broader field of computer science, wasn't it?

Landwehr: I was trying to; it was educational for me, but also it's when you try to write it down then you learn what you don't know. So it was an opportunity for me to be in touch with a broad range of people doing work in that particular area and to try to write it very carefully and precisely. That was one of the other things that was around the communications system staff at the time was, as soon as you came onboard you got a copy of Strunk and White [*Elements of Style*]. I knew Strunk and White before that, but I have to say, I think that working in that environment and having people review internally,

and think about that, made me very careful with what I wrote down. I remember, Dave Parnas said something at one time about he had never written a paper that he wanted to go back and fix because he's very careful before he put them out.

Yost: Focus on clarity, being concise, and getting things right.

Landwehr: Yes, exactly. But I thought that was really good, a good standard to hold yourself to.

Yost: In 1970, in Willis Ware's Defense Science Board report, one of the things that he emphasized at the end of it was the importance of industry partnering with government to find solutions. Did you give any thought with writing that *Computing Surveys* paper that this would also educate researchers and scientists in the industry to do more with developing secure systems?

Landwehr: Absolutely. I viewed it as an educational tool really for a broad community.

Yost: I see that you presented on multi-level security at the Formal Verification Workshop in April 1980. Can you talk about the workshop and what you feel was achieved, and also, the reception to your paper?

Landwehr: I was reading about some of the work that was going on, and the position I was in, I felt that there were people out there working on developing verification systems



and verifying properties, but they didn't necessarily want to know what properties they wanted to verify. I thought it was an appropriate role for the government side to be able to try to specify what are the properties that we want in these systems. So as I recall, anyway, that there was a verification workshop, first there was a workshop. I think that was the first time I was over at SRI I don't think I had interviewed there. And I remember I was very happy to be able to go, because I thought this was a group of people whose work I had looked at and respected, so I was glad to be part of it.

Yost: In the early 1970s, Roger Schell led a program with the Air Force that supported the two Anderson Committee and the development of its two reports, as well as work at MITRE and Bell-LaPadula among other projects. When were you first aware of this research program of the Air Force, and what were your early impressions?

Landwehr: Oh boy. I think I was actually aware of that; I can't say it was at Purdue, for sure, but I think when I was working at the Navy Yard, because there was discussion about security kernels at that time, you know, within the military side trying to do multi-level security. So I think that's probably where I first looked at it. And then when I was doing the security modeling work, I looked very hard at Jerry Popek's work, his dissertation. There were a couple of SOSP meetings were held out at Asilomar [Conference Center, Pacific Grove, CA] around that time. I remember that's where I met Dick Kemmerer the first time, I think, out there.

Yost: That acronym?

Landwehr: Sorry, it's the Symposium on Operating Systems Principles. So, I think there were two consecutive ones were held at Asilomar and I think that was after I was at NRL, otherwise I wouldn't have gotten there. Dick was at that one and I remember when we did our work on military message system security model, we looked very closely at Popek's work and Dick's work on what are the information flows in CPU instructions.

Yost: I know prior to starting Trusted Information Systems. Steven Walker in various leadership positions in the federal government played kind of an organizing role and specifically some NBS/NIST workshops. Were those ones you participated in, in the early 1980s?

Landwehr: Yes. When I first met Steve, he was still at DARPA, he had left NSA. He was sponsoring the SCOMP work and I think he and Jerry Popek were at some of those meetings at Honeywell at St. Petersburg. And then when he moved to the Pentagon he started this computer security initiative. I can't remember what year that was; it was probably 1978 or 1979, I don't remember. And so I used to get invitations from him to go to meetings at the Pentagon, and generally, I went to those. The first Bureau of Standards meeting, I'm not sure I was at that. I don't remember; but I went to most of them after that.

Yost: Was there more a sense of cooperation or was there also a degree of rivalry between the Air Force researchers and naval researchers with computer security, in the early years?

Landwehr: Well, I think it was pretty cooperative, really, overall. Certainly, the computer security initiatives were cooperative among the services. Researchers are always competitive, to a degree. What I think I can remember — and this is a long time, now — that Dave Parnas was skeptical about the value of security kernels and actually Frank Manola also said; really, the things that seemed to work were not; this centralized checking mechanism was going to be too inefficient and it just wasn't going to work. And so, okay, maybe that's a possibility and so there were prototypes being built, [and] did they work or not seemed like questions you could answer. But there was also the question about the Bell-LaPadula security model, and was that really going to be practical or not? Again, DARPA funded prototypes, SCOMP and KSOS, and the experience with those was interesting in that when you tried to build things based on that model, you ended up with a fairly large category of trusted processes, things that were privileged to violate some rules of the model. But you hopefully had some trust in them that they weren't going to do that [misuse the privilege]. And you actually had to violate those rules in order to meet the functional needs of the system. And that's when we started working with the military message system model, we wanted to try to — never mind the architecture of guts of the system — identify what are the functional requirements and what are the security requirements on the functions. So we tried to specify those together.

Yost: That's when you started to look at application-specific [pause]

Landwehr: Yes. That was an outgrowth, and Connie Heitmeyer was a very important player in all this also at that time. The military message experiment was this activity in which PDP-10 was put into CINCPAC. At the time, CINCPAC had a well-established pneumatic tube messaging system, as I understand; I never saw it but I believe the Mayo Clinic had such a system that was still in operation in the 1990s, because they could ship around medical records that way very efficiently. The Navy was doing the same thing with military messages at the time. DARPA was pushing time-sharing systems, TENEX and so on, the Navy had these LDMX and NAVCOMPARS systems, which were really descendants of sort of telex type of messages, but the Navy had been using these kinds of messaging for a long time. So the military message experiment at the time, was to try to put in a PDP-10 running TENEX on top of a simulated security kernel — I'm sure Steve Lipner probably told you all about this — to see how it would work and in fact, to observe the operational use, and what the users did. In order to comply with the artifice that there was a security kernel underneath, they wanted to have the users approving operations where the security model underneath was being violated by something. The key finding out of that was that users had no idea what they were approving, and it was, at the time, the equivalent of a dialog box that would pop up and say, "do you approve this?" And the users would say yes, and then afterwards, if you interviewed them and asked them why, they would say "because I want to get my job done"; not that they had any concept of what was going on. So Connie was a witness for that stuff and wrote some of the reports about that.

Yost: What year was this?

Landwehr: That's a good question. I have some of that stuff upstairs. I think that was going on probably in the late 1970s, between 1976 and 1980; probably 1976-78, I'm thinking. One of the guys from NRL, John Kallendar, really wanted to wind up in Hawaii and he managed to do that. But, in fact, what happened to the system in the experiment was, no one really thought they'd send the PDP-10 back, but they did. I think it was a change in command or something like that. A new guy came in and said I don't want this stuff around. So it didn't succeed and I think DARPA was trying to do technology transition by plugging that in. Anyway. So that work, what I'm remembering which office I was in when we're talking about this stuff. I think it was in Building 54, which was the first building we were in when I was there, and so I'm thinking it was before 1980.

Yost: When did you first start attending the National Computer Security Conference and what were your impressions of that conference in the early years?

Landwehr: That year was really, as I remember it, the descendant of those NIST or NBS conferences, and so I started attending those pretty early on. As I say, I'm not sure if I was at the first one, but I think I was at most of the subsequent ones until it became the National Computer Security Conference [NCSC], and then it got to be a convention, really.

Yost: Quite large.

Landwehr: Yes indeed. And NSA put a lot of manpower into it and its funding, and in the end, they decided to stop doing that. I'm not sure how the decision was made. I think the last one was around 2000; it was after I'd gone to Mitretek. It was always kind of a mixed kind of conference. It was not a high level research conference, it was a place where you could go talk to people and maybe hear about experiences and see what technology was coming down the road, what the defense community was producing.

Yost: When did you first attend the IEEE Symposium on Security and Privacy?

Landwehr: Again, that was started as a workshop one year, and I was not at the first one, but at the second; so it was in 1979, maybe, was the first one? Or 1980?

Yost: I think the first workshop was 1980, or organized in 1980 with 1981 being the first symposium.

Landwehr: If so, then 1981 was the first one I attended. KSOS was being developed by Ford Aerospace, out in Silicon Valley, and I was one of the people on a review team for that, and then there were a bunch from MITRE. And I think the MITRE people set up the first one and the motive was to have it in the spring in California, where you might be able to go skiing in the spring. [Laughs.] I heard about this from one of my grad school

colleagues, Jim Hamilton, who said are you going to this? I said no, I didn't hear about this. MITRE guys had set it up, and there was some competition between MITRE and NRL, though in general, it was pretty collegial. But the second year, I was invited.

Yost: Can you talk about that event, and as I understand, there was much more of a scholarly rigor to it than the NBS/NIST/NCSC events?

Landwehr: Yes, in the first couple years, actually, they produced a pre-proceeding. It was a much smaller community, of course, at that time and the reviewing was done, I think everybody read everything, as I recall, and there would be some discussion about what to present. And actually, authors had the opportunity to revise it following the meeting because there would be vigorous discussion at the meetings and the proceedings would come out after the meeting, which seemed like a good thing, actually. I was involved in that pretty much from the start. There was normally a dinner the night before. The meeting was normally Monday through Wednesday noon, as it still is, I think, and there was often a dinner the night before, which was really more the organizers got together and decided who was going to be program chair the next time, or whatever. I was invited into those meetings and participated in those, and got involved in organizational things that way.

Yost: You talked a bit about the *Computing Surveys* article. How was that received by the computer security research community and did that help establish yourself as a central figure in the field?

Landwehr: I think it was important and I think people read it, and then they knew who I was, and so that helped, for sure, I think. I got involved in something that's maybe less publicly known, but there's a treaty among Australia, U.K., U.S., New Zealand, and Canada, the TTCP which is The Technical Cooperation Program, which is an instrument for enabling defense labs in those countries to exchange research results. And I really got into that, I think, because Marv Schaefer, who was a chief scientist at the National Computer Security Center at NSA, he got me involved. And I think also in the early Woods Hole Study, as well. The fact that I had written for that, and I was in government, there weren't very many people in the government who had written papers that were in *Computing Surveys*, so that helped.

Yost: In the mid-1970s, IBM SHARE first started having sessions on computer security and a product emerged from engaging customers, RACF, which down the line tried to qualify, and only qualified for C1 [in TCSEC]. Was it frustrating that industry wasn't doing more in the late 1970s and 1980s, especially IBM with all its resources?

Landwehr: Yes. And it's still frustrating; I'm still thinking about how to make that sort of thing happen. Industry has, obviously, a profit motive and they have to spend their resources where they think profit can be maximized. And of course, the government was trying to make it satisfy what it felt its needs were for secure systems, while working with the industrial incentives. And that's sort of what the whole Trusted Computer Security



Evaluation Criteria was about. It didn't work out as people had hoped. But yes, it was frustrating that we couldn't get more out of industry than we did.

Yost: Were there any individuals in the computer security community that were writing papers on the economics of computer security in the early years? I searched a bit and I haven't found much of anything.

Landwehr: I don't think so. At least not that I was aware of, or am aware of, I guess. What I remember very well was the call for papers for the first workshop on the economics of information security; that's Ross Anderson and the economist at Berkeley, what's his name? He's now at Google. I'll think of it.

Yost: Is that Hal Varian?

Landwehr: Yes, Hal Varian, I think Ross and Hal were there, and maybe some other people, but I know they were both there. And I had thought, this is a workshop I want to go to and I figured out how to write something that would cause me to get there.

[Laughs.] And I did. Trying to think; there was an early workshop that was on integrity protection that was held at a business school outside of Boston, much earlier on. I think after the Clark-Wilson stuff came out, so there was some sort of business interest there.

And, of course, in the early days, there were business guys. There was Bill Murray and Bob Abbott, and who's the other one I'm trying to think of who is also passed away but was [pause]

Yost: Bob Courtney?

Landwehr: Yes, Bob Courtney. They were very much in evidence at a lot of functions and generally speaking, felt that what industry was doing was just exactly the right equation. That this work was well justified, put it that way; there was no need to worry about these other things. And maybe, you know, at that time, the threats weren't there. Now I think it's pretty clear the threats are there.

Yost: In the 1981 *Computing Surveys* article you singled out two industries or economic sectors, finance and medicine, where there were laws that provided particular challenges for designers of models. Can you expand on that a little bit?

Landwehr: I did? [Laughs.] This is the formal models paper?

Yost: Yes.

Landwehr: Oh, well good for me. Boy, I need to...probably got a copy here somewhere. Here we go, no. So that was 1980, 1981. Here we go. Where was I talking about [pause]

Yost: I think it's late in the article, that these sectors were presenting challenges, based on existing laws.

Landwehr: Interesting. Okay, [reading a selection], yes okay. Well, that's certainly true. I don't know that I; last sentence of the paper. One thing actually that I should get credit for in this paper that's probably not anywhere else, is the fact that the *Computing Surveys* had a very good editorial process, so not only did I work hard to make this understandable but I got good feedback from the editorial process, which included it, also. But I don't remember what was in my mind at that time about those systems, but certainly later on, when I started the working group and I wanted to find examples that were not military in which security was an issue, the medical example was front and center. Financial I probably haven't explored as much. I think the other thing that probably brought that up is if you think about aggregation, things like that, those kinds of examples tend to show up.

Yost: In 1984, you published an article in ACM Transactions on computer systems in which you outlined four major systems — MME, AFDFC, Multics, and KSOS — where you found the Bell-LaPadula model overly restrictive in practice. We talked a bit about that but can you expand upon that a bit, and also talk about exceptions, and about this critique of Bell-LaPadula, at that time?

Landwehr: I think I have touched on that, as to why. We were really trying to think about including security as a fundamental property in the system, rather than thinking that you could actually just secure this bit and the rest could be whatever you wanted. So we were trying to think in terms of security as a function. That was one of the discussions, is

security a functional requirement or a nonfunctional requirement? Lots of people now will even tell you security is a nonfunctional requirement. But in some sense if it has no effect on function, then it's not a requirement. And that was, I think, the view that we were taking. So we tried to take functional specifications, which the software engineering group at NRL, and Connie in particular, worked on specifying family members, trying to follow Parnas' family methodology for specifying systems. So you have more constrained members, and more capable members, and there may be overlapping subsets, or nested subsets, of different arrangements. But the question was, how do you bring security into those specifications, and we thought the way to do it was really to try to do a rigorous job of specifying the security within the framework of those functions. And so that's what we really worked on. The reception, I think it was; well, the thing at the time that seemed novel and that people felt typically was lacking from the Bell-LaPadula model was multi-level objects. So objects in which you've got a security level attached to something, but it may have something inside with a different security level attached to it; there's only a lower one. And that seemed to be a fairly intuitive kind of structure to have in the model because it matched, you know, we have a TS certified safe, but it may have confidential folders in it. You take the folder out, it's okay, it's not polluted by the TS safe. And that was a problem with the Bell-LaPadula model that people perceived and so there was a lot of discussion, I think, about multi-level objects. There was a complaint that we didn't provide this kind of kernel structure, this boundary that a lot of people had come to think of as the right way to do it. Including Steve Lipner, actually, at the time, he complained. I think there's a *Computing Reviews* review that perhaps Steve wrote — not sure — but complains about that.

Yost: And this was at the time when he was at DEC trying to build a kernelized AI system?

Landwehr: Yes. I think it was before that, probably. It was in the early 1980s, right?

Yost: I think that probably started about 1981.

Landwehr: I'm not sure when they started the SVS work.

Yost: I think it was about 1981 to 1988 or 1989.

Landwehr: Okay, then it would've been. Yes. So we went ahead and built prototypes based on that demonstrated function, but we didn't try to build one that we felt we could actually certify. That was something we felt industry ought to do for us, but we weren't contracting for that.

Yost: Was the MMS the first role-based access control?

Landwehr: I guess so, in my view, I don't know of any earlier ones. And we did very definitely talk about that concept because when you think about the functions and message systems, you've got people who stand watch, you have people who release messages. So you may draft messages but somebody else has to be the releaser, and it

can't be you. So you very definitely have this notion of different roles and different privileges associated with the roles. And you might be allowed to become a releaser, maybe not for your own message, but you had to think about these relationships and so we included structures in the system to deal with that, which I think at the time we called access sets. We did talk about roles because what I've just described is a role. I remember when the Clark Wilson stuff came out I thought that's exactly what we were talking about but nobody had seemed to notice.

Yost: I understand that the 1984 paper formed the basis for the National Academy of Science's sponsored study of database security problems in the 1980s. Can you discuss that and the key individuals who made meaningful contributions?

Landwehr: This is the Air Force study that we referred to for a long time as the Woods Hole Report, although there were lots of Woods Hole reports by then. So this is a multi-level secure database study, I think Marv Schaefer and I don't know if it was Dorothy Denning who co-chaired that? She was certainly involved. Things have changed a lot since then. There were three sort of tracks in that study. One was intended to be the near term, and that, as I remember, Roger Schell had this sort of what we called a spray painting approach that dominated. And there was sort of medium term, which is where the MMS work fit, which was viewing the message system as a database and sort of going upwards databases from that. And then there's the long term research group that Dorothy headed, and which I spent a good bit of time on; studies don't seem to do this so much anymore, but at that time, there was like three weeks at Woods Hole and I ended up

being there the whole time. It was very stimulating and that's actually when I think — non-interference — Joe Goguen and Jose Meseguer - put out a non-interference sort of model for that structure. So there was a lot of discussion and then there was writing up of reports after, which took a long time to get all that reviewed and released, which often happens with academy reports, I guess, but I wasn't so aware of it at that time.

Contributors: Dave Bunyan from Canada was there; Connie was there, Connie Heitmeyer; and I was there; and I guess I'm not sure if John McLean was up there or not. You can find out tomorrow. I don't think so. I'm not sure when he got to NRL.

Yost: Of course, you and Connie and John were focused on computer security, were there other people focused at NRL on computer security as well?

Landwehr: Well, there were people in my group, certainly, because Mark Cornwell did a lot of the work on prototypes, and there were some people who came and went; some of them came back and are there now. I'm afraid I'll perform injustices in not getting them all named, but they're mostly co-authors on various reports, and so on. You know, I had a group that varied in size from probably four to eight. But Judy Froscher was an important person later on, and recruited people, and had people also, sort of as a continuing group now.

Yost: In 1983, you published an article, "The Best Available Technologies for Computer Security" in *IEEE Computer*, which like the 1981 article, brought the appeal to a broader

audience. Was that something that you decided you wanted to continue to do? Or did someone at *Computer*, did they approach you?

Landwehr: Nobody approached me about it, no. I don't know if I'd seen things in *Computer*, or maybe someone may have said there was going to be a special issue on security or something, I think that probably triggered it to some degree. But also, I can remember having conversations with Dave Parnas where he complained about there were all these different things, that people were defining things in different ways and there was a need for something to bring that together. And so I attempted, along with the descriptions, to try to provide a set of definitions for things.

Yost: Terminology. One of your last summarizing points of that article is there's lots of "modern ideas that don't work yet." So in 1983, what specifically were you thinking of as modern ideas?

Landwehr: I think that was a dig at security kernels.

Yost: Okay.

Landwehr: Maybe capabilities too, but I'm not sure about that.

Yost: Could you speak a bit about your view of capability models?



Landwehr: They always seemed appealing to me. They seemed like the logical structure matches the problem, and what I learned only relatively recently was that, in fact, they've been hidden inside of successful IBM systems for quite a while. And so that's, in a way, a reassuring thing. I learned that from, I think, conversations I had with Paul Karger not too long before his untimely death, when I was working on that summary paper looking back over the history of funding for security research.

Yost: You also mention that auditing faces the dilemma of being too voluminous a task to completely do manually, yet automated systems are doing it in a way that becomes kind of useless. Is there a middle ground there? Is there a kind of appropriate level...

Landwehr: You mean how much to audit?

Yost: Yes.

Landwehr: I think that's still a challenge. You think that you want to audit at the level at which the transaction makes sense. A lot of auditing is very low level. On the other hand, if you end up having to do forensics, it may be that that's what you need and the diminishing cost of storage certainly has made it possible to collect a lot of stuff, if you want to. But it's not clear that it makes sense to collect it if you're never going to use it. That I think is still a challenge.

Yost: I believe it was in 1980 that James Anderson first highlighted the idea of intrusion detection, as well as automated systems associated with intrusion detection. And, of course, still early in that decade, the SRI IDES project starts. What was your view of this field of intrusion detection expert system research as it was getting underway?

Landwehr: Actually, it's interesting. I heard people talk about that sort of thing, that is trying to figure out whether there was an intrusion going on because you had an abnormal pattern of access to a system, like somebody logs in at the wrong time. Years before, I heard Stan Wilson talk about it, and H. O. Lubbes, also. And I think H. O. may have actually funded the IDES/NIDES work, or at least some portions of it at SRI. So I think that idea, I don't know if it comes from Jim Anderson or not, but it's a pretty straightforward idea, in that sense. From a personal perspective, my predilection is rather keep them out than worry about detecting them after they're in. But certainly, there's no question that we have to do that.

Yost: Was Anderson unusual in that he was very interested in kind of both sides, whereas many focused on high assurance have less interest in intrusion detection?

Landwehr: I think Jim was unusual in many ways.

Yost: Dorothy Denning has been interested in both sides.

Landwehr: Yes. Jim was just a great and a very helpful guy. He had interests and exposure to such a wide range of problems and interests that it has always been good to talk to him.

Yost: There's some other areas that you worked on directly or oversaw the work of other scientists at NRL that I'd like see if you have some comments, like the NRL pump.

Landwehr: [Laughs.] So the pump was not my idea, but I think I may have come up with the name. I think Ira Moskowitz and Myong Kang deserve the lion's share of credit for that. I think there's something that; let's see; simple question, I don't know when the first pump papers appeared, but I think it's in the 1990s, right? I was just looking through these, trying to get the first one.

Yost: That sounds right.

Landwehr: Something that I mentioned before about TTCP, this Technical Exchange, Technical Cooperation Program, it's interesting. The same year that I started really getting involved in that, I started the IFIP working group on database security. I always thought that the IFIP group would be more productive, in some sense, and the TTCP work was just sort of because we were part of the Armed Forces and we were supposed to do these things. But it wasn't so simple. I think the IFIP group has been productive in its way; it certainly has produced lots of papers. It hasn't taken the role that I hoped it would, exactly, but that isn't what I wanted to focus on. TTCP, what was interesting, was

to be exposed to the work that was funded by smaller countries, if you will, in particular Canada, but to a greater degree, Australia. I found the Australians often had different ways of approaching problems because they knew they weren't going to have a big effect on industry, so they had to figure out what our industry was producing. There was a woman, actually, well, there were a couple of people; Brian Billard initially, and then Mark Anderson, who has just recently, I think, retired down there. They had very innovative ideas of how to do things with what you might refer to as gadgets, and the pump is a sort of a gadget in that sense. The Australians had produced something; well, people had built one-way flow devices, and they still do, out of simple fiber optic connections of one sort or another. But they didn't allow for the protocols to work, and so the innovation with the pump was to provide a back channel that you could monitor that would be just enough so that you could make the protocols work but you wouldn't provide covert channels that would be intolerable. That's been quite successful, I think.

Yost: When did you first see covert channels as an important problem with computer security?

Landwehr: They've been around for a long time. [Laughs.] I can remember [pause]

Yost: I got the impression that the vulnerabilities just continue to escalate, and that people knew about the potential vulnerabilities that could be exploited.

Landwehr: Yes. I can remember certainly people talking about, you know, at the time, driving a teletype at 300 bits per second, and off of a covert channel. This must have been early 1980s or late 1970s when a teletype wasn't such a strange thing to have and driving it at that speed seemed like that was full speed. [Laughs.] And off of covert channels, you know, so people have known they were there, but felt they were not the main threat for a long time. Then, I think, there's still a degree of that.

Yost: I remember Steve Lipner emphasizes that as one of the great challenges.

Landwehr: Yes, they took it very seriously, I think, especially with Paul Karger on the job.

Yost: What about onion routing?

Landwehr: So that's another one I can't claim any credit for, but I was around when it was being developed and I certainly had many lunchtime conversations with David Goldschlag and Paul Syverson, and Mike Reed about that. I think I may have helped recruit Mike Reed into the organization and I did a little bit of analysis on it, with Gene Tsudik; to try to figure out what security was really providing or what the routing provided you. So I was very interested in it and wanted to see it developed, and I don't know, it's another one of those things that's turned out to be a good thing for some and a bad thing for others.

Yost: A taxonomy of security flaws?

Landwehr: So that really grew out of, I think, in the time period around; I started accumulating more, in a way, collecting the gray literature again and trying to bring it to light because there started to be the accumulation of these kinds of exploits. But there actually had been some written about a long time before, by Bisbey and Hollingworth; in fact I think there's early papers in the PAP [Protection Analysis Project]. So there had been some history of that, and I thought the information we were collecting, appropriately organized and described, would be valuable to others. And then I was confronted with how to organize it, and I was going back to the software engineering side of things and saying what I think what would be useful is how to avoid introducing these flaws into your system, so I developed a scheme for organizing the data that I thought would be helpful in that way. The frustrating part of that experience was that it was very hard to get anybody to openly admit they ever actually organized their data that way and took any action based on the result. It's possible that it was used internally in some companies; I wanted to get the Software Engineering Institute, the SEI/CERT [Computer Emergency Response Team], to use it to organize their data. I wrote to them, I wrote to their sponsors, nothing happened. I even complained about it but what happened instead, it seemed, that this was a good topic, creating taxonomies, and people wrote dissertations on different sorts of taxonomies as a result, often attack taxonomies rather than flaw taxonomies of various kinds.

Yost: We discussed James Anderson as an early key contributor. I wanted to get your impression of Ted Glaser, did you get a chance to meet him?

Landwehr: I didn't know Ted very well at all, but I think he was at the Air Force Summer Study, the Woods Hole study, and I think I met him there. He was clearly a perceptive guy but I had not much interaction with him before or after, really.

Yost: In 1983, you became a supervisory computer scientist. Can you talk about the difference and the transition of that in your career, from being more directly focused on research versus research and also supervising?

Landwehr: It was a promotion. I don't remember what promotion it was at that point, but it meant that I did have to do the performance reports for people. But NRL was a pretty flat sort of organization, really, and so people generally worked pretty independently. I don't remember it as a huge change of any sort, really.

Yost: Talking to Peter Neumann, I got the impression that SRI is kind of similar.

Landwehr: I suspect so. It's really a research lab and the idea is don't hire people that require a lot of supervision because you won't want to do that. [Laughs.] That was one of my other successes in that respect, was I competed with what was then the Bureau of Standards, I guess, probably Denny Branstad to hire Cathy Meadows, and we persuaded her to come to NRL. That was a success, but I don't think I ever supervised her, actually.

Yost: Can you discuss her, then, in terms of her contributions?

Landwehr: Yes, she's a brilliant person who has done a lot of innovative research over the years, particularly, of course, in the area of cryptographic protocol analysis, but she was always a good person to talk to about whatever technical topic. Same with Paul Syverson, whom I didn't recruit but I think John McLean did. I think John was one of; John we got from North Carolina probably with Dave's assistance. I think leaving NRL was overall definitely a good thing for me, but it certainly made me appreciate the environment there, having left it.

Yost: How did you see the evolution of the infrastructure for computer security research, in light of your two decades there?

Landwehr: The evolution of the infrastructure.

Yost: Was there significantly more funding?

Landwehr: I guess it had its ups and downs. The 6.1 funding typically went to people like Paul and Cathy, who were doing more logical mathematical stuff, and the outside funding did fluctuate. Well, the particular thing I remember and I think I probably put it in that paper was the BLACKER Program at NSA, at one point, seemed to soak up all the research dollars. The security initiative, the idea was to try to coordinate what the



different services did was with Steve Walker's thing; not really to centralize the funding, I would say, but to coordinate the funding. But of course, they created the National Computer Security Center then suddenly the funding was all gonna go through there. But the plan was that it would all then be redistributed in a consensual sort of a way, a collegial sort of a way, to the services and I think that happened for a while, but then when NSA ran into trouble on something and needed that money, then suddenly the services found that they were going to have to fund it through their own budgets again. It varied, but I don't remember that being a big problem. The worst problem was, and what I think ultimately contributed to my departure was I felt that we couldn't recruit people into the salary structure that the government was providing, particularly when people were all going to participate in the Internet bubble, or what turned out to be the Internet bubble on the West coast. It was just we recruit somebody good and then lose them because of that sort of competition and that makes it difficult to survive.

Yost: You mention the National Computer Security Center. What were your thoughts when you first learned that there would be such a center and that it would be a DoD center housed at the NSA?

Landwehr: Of course I knew about it, you know, because the computer security initiative and Steve Walker talked about it. The rationale for its creation was really evaluation and I think I talked to him about would I want to be part of this? I said no, doing routine evaluations doesn't sound like something that I wanted to sign up for, but I could see the rationale for it. I know he actually wanted it to have it at, I guess it was

probably still NBS at that point, at the Bureau of Standards, not at NSA and there's a story there. You probably know about that as to why it ended up at NSA. But the thing that I remember thinking at the time was it was really created as the National Computer Security Evaluation Center, and they immediately dropped the "Evaluation" from the title and made it the National Computer Security Center. And I think those of us who felt we were doing related research at NRL, felt to some extent this was a land grab. But I think, there's nothing in particular you could do about it, number one; and the only thing you could really do about it was do good work and see that it got funded; and we were generally successful with that.

Yost: At the time, were you optimistic or more pessimistic about industry response to setting up this criteria and evaluation infrastructure?

Landwehr: I think I was probably more optimistic, at that point. I mean, I thought the criteria made some sense. I remember discussions at the time — from maybe Denny Branstad, or maybe it was others, I don't know — that it made into a standard. I thought no, it wasn't written as a standard and at the time it was inappropriate to do that, yet that opinion held the day. It seemed as though companies were signing up; they were gonna build things and have them evaluated, and we were going to get somewhere.

Yost: Did you feel the structuring of the levels made sense or did you think there was a better way of doing it?

Landwehr: Naturally, we thought there was a better way of doing it and wrote things about that, to some degree. The difficulty with that was that by putting assurance levels together with the functions you overconstrain the problem. That turned out to be true, particularly as we started to get components, it didn't make sense; it was really written with this model of an integrated system, a time-sharing system. And that's not, you know, the direction the future took for computing, and so then trying to evaluate particular components against these meant that some of the things didn't really make sense. So we were more in favor of what Roger Schell would've described, I suppose, as a Chinese menu approach where you'd say what the functions were and what the assurance was. And that, in the end, is sort of what we've ended up with the Common Criteria, and I can't say that that's been a big success either. So it's a challenging problem and I don't think the solution has been discovered yet.

Yost: If I understand Roger Schell's reasoning correctly, it was that the structure that it took would be easier for defense contracting with fewer categories, as opposed to the Chinese menu?

Landwehr: Yes, and maybe that was the right thing. From a technical standpoint it didn't seem like the right thing; and from a contracting standpoint, you know, it certainly constrains the decisions that contractors just have to make, and so it's simpler in that way.

Yost: I understand that you were involved with the formation of the IEEE Computer Society Security and Privacy Technical Committee, is that right?

Landwehr: I think it was formed without my help.

Yost: You were involved early on?

Landwehr: I was involved early on. The Technical Committee, I think it really happened because they needed something off of which to hang the Oakland Conference. So there needed to be a home for the conference in the IEEE hierarchy, and it may have been Steve Lipner, I don't know; Stan Ames was another senior MITRE person at that time, it may have been Stan, or Pete Tasker was someone else; I really don't know who was the first, who did whatever it is you have to do in the IEEE to create a new technical committee but it wasn't me that created it. I've got records of this, but there was an initial — I'm not sure what's the right title — whatever, TC chair. Yes, chair for the technical committee, TC chair would go to the TAB meetings, the Technical Advisory Board meetings at IEEE, and the initial TC chair, I forget who it was. Then I became the vice chair. And that was usually, you spent two years as vice chair and then spent two years as chair, and I did that but I was probably the third TC chair, I'm guessing. Because I think Dick Kemmerer was chair before me, if I remember correctly, and somebody was before him. But I'd have to look up who that was.

Yost: Did it evolve to have any particular roles that weren't connected to the Oakland Conference?

Landwehr: Not a lot. It became a home for some other conferences that the TC; and I think there was the Foundations Workshop, which is now maybe a symposium that started with IEEE sponsorship when I was there. The new paradigm, I don't know; it was funny, the real competition at the time as I recall it, was the ACM. The ACM didn't have any security special interest group. That's not true, they had SIGSAC pretty early. But SIGSAC tried to organize a conference, and I remember actually we sent a paper into this from NRL at the time. And in the end, I think they didn't get enough submissions or something, but they never held the conference. The IEEE was sort of the only game in town for professional society security stuff for quite a while until the ACM CCS Conference got started, whenever that was. But it was quite a while later. Of course now, it's been [pause]

Yost: Any reasons you see for that?

Landwehr: I think that maybe academics were less interested. I think of ACM as having a bit stronger interest in the computer science side of things; well, definitely stronger in the computer sciences. IEEE has got a lot of exposure to engineers and, I think, a little bit more toward the defense side of stuff, and so it may have drawn out that way.

Yost: And were you involved with the founding of *IEEE Security and Privacy*, the CS magazine?

Landwehr: The magazine? I was, but George Cybenko really was the prime mover for that and you should talk to him about that. But he recruited me and a few others to the editorial board as associate editors in chief, or something like that, early on. That was when I was already at the National Science Foundation, I think, before that happened.

Yost: And what were some of the founding goals for that publication that you remember?

Landwehr: I think it was to satisfy what he felt was an unmet need for a regular publication in that area. And I remember at the time, it was controversial. Some people felt that the last thing we needed was yet another publication, but I think it's turned out that he was right.

Yost: The *Journal of Computer Security* existed at that time.

Landwehr: Yes, the *Journal of Computer Security* was there [and] that wasn't the only one either. There was *IEEE Transactions on Dependable and Secure Computing*, I guess, I'm not sure if *Dependable and Secure Computing* predated the magazine; they were pretty close to the same time. I had some involvement with that also because through my

relationship with IFIP Working Group 10.4, so I knew that was going on and I was on the editorial board of that, initially, also.

Yost: So in 1999, you left NRL to become acting director of Mitretek CIS? Can you talk about that?

Landwehr: Well, I didn't actually leave to do that, I left to become a Senior Fellow at Mitretek. So I was recruited, basically, there was a headhunter out there looking for people for Mitretek. I think John McLean was recruited at the same time but he turned them down. I was frustrated because okay, I'd spent many years now at NRL, and I could easily spend 10 more and retire, but it didn't seem like we were making any progress in the real world and we were having a hard time retaining new hires. And of course, there was a substantial increase in salary. I thought maybe this is a good thing to do, and maybe I can have more effect on real systems. And when I got there, of course, Steve Lipner had recruited me into that position along with Stone — what was her first name? — there was a vice president who was a woman named Linda Stone. Recruited me and I thought okay, this will be interesting. The idea was that the senior fellows, and they had recruited a couple of them, including one that they got from the CIA, were going to provide senior oversight and consulting on their projects. The way it worked out was not too long after I got there, I took a vacation, which I came back from vacation to discover that Steve had taken a job with Microsoft and there I was. That was okay. I hadn't really counted on that but we'll keep going. But also, the company had lost some relatively long term and lucrative contracts that they had with the Post Office or somebody unrelated to

what I was doing. But it meant that they suddenly felt they had to have everybody fully covered on a contract and so they basically they said, you know, you need to find somebody who's willing to hire you. Happily, Sami Saydjari at DARPA at that time was running a lot of information assurance programs and new ones with other PMs and they needed help. He was enthusiastic about it so happily, I got support. But what it meant was I spent my time helping DARPA, which was fine, and I don't mind having done that at all, but I didn't feel I was doing for the company what I thought I had been hired to do for the company. And as a relatively well paid person there, the incentives were set up within the company so that if somebody wanted you to consult on something, they were going to have to absorb your salary or some fraction of it. And so I really felt that what I learned there was what it would be like to live as a lawyer and charge every hour to somebody. That wasn't, you know, I could do it but that wasn't what I thought I had signed up for. So then when I was approached by NSF to be an IPA there and start running programs there, that sounded attractive and all I had to do was to get Mitretek to agree to it. And they did. In the interim I did spend time as acting director of that organization, then they hired somebody, John Davis, an excellent person, whom I'm still friends with, and he was happy to let me go to NSF. NSF was going to pay my salary and that was fine.

Yost: What can you tell me about NSF developing a research program in computer security? Who was the prime driver or who were the prime drivers?

Landwehr: The program — and I want to give them credit because really, it was created before I got there, in the sense that they already had a solicitation out. I believe that the



prime movers on that were Helen Gill and Kamal Abdali, who was the head of that group at NSF at the time. George Strawn may have been acting CISE director at that point, I'm not sure. I kind of think that was the case, and after I'd been there about a year, Peter Freeman came in as CISE director. And he certainly deserves credit for building the program substantially and then subsequently, other division directors under him helped; Greg Andrews and Wei Zhao and Taieb Znati, subsequently, after that and now Keith Marzullo. So a lot of people, but I think Helen Gill, I suspect, drafted the initial solicitation or contributed strongly to it.

Yost: In stepping into that role, what were your goals with the program?

Landwehr: To tell you the truth, I was happy to be there and be part of the process. I had done some reviewing for NSF when I had been back at NRL, but that was more in the time when they — at least for the programs I was involved in — were more often doing mail reviews than panels. I thought this was a great opportunity for me because a lot of times, a typical NSF rotator comes in and they inherit a program that's already got a lot of things going and maybe they got a little new money that they could spend on a few new proposals, but they wouldn't have a lot of latitude. I came in with not a huge amount of money, but basically, uncommitted money and I could start the thing from scratch. So my goal was to get the research going that would have some effect and it was a great opportunity.

Yost: It's probably every program officer's role or practice to lobby for more funding for their program, but I came across an editorial of yours in 2005 citing that there were 390 proposals — actually, that was for the year 2004 — but there were only funds for 35 awarded projects, less than 10 percent funding rate. We had another project on the history of FastLane so got some sense on funding rates agency-wide. I understand that funding was more on the order of 20 percent for most programs, or at least, on average. Did you feel that the funding just wasn't adequate in your first tenure working at NSF?

Landwehr: I guess the advice from George Strawn — he was asking me how many proposals I got and how many I was able to fund, and so on — he pointed out that it was a good thing that there was proposal pressure if you wanted your program to grow, and it was no problem to exhibit that. But I can't say that I felt that funding was at such a level that I was forced to pass up things that I really wanted to fund. I mean, you get a lot of proposals but they're not all great. And it's often a process whereby a proposal comes in and it gets turned down with some good comments, and the proposer can come back the following year with a better version of it; maybe that one will be funded. So, no, I don't think I felt that it was; it was clear; it's frustrating, in a way, from some perspectives.

We've got lots of researchers in the community, there are lots of jobs in cyber-security, but what they indicate is that what we started trying to do in the 1980s and the late 1970s hasn't succeeded. We've got huge problems and so people are willing to spend money on it because they can see we have problems. But whether we're solving those problems this way is not so clear.

Yost: You were on assignment as a program officer, much of the first decade of the new millennium, but in 2003 you also picked up an appointment at University of Maryland as a senior research scientist in the Institute for Systems Research. Can you talk about that institute a bit, and the role that you had there?

Landwehr: First, the thing to understand is what happened really, was that after I'd been at NSF for a year, NSF said they'd be happy for me to stay for two more years, and my management at Mitretek, my immediate management, was quite happy for that to continue, but the CEO of Mitretek was not; she wanted me to come back. I said well, we've just put out a solicitation, it wouldn't be fair at this point to leave. And so she said pick a number from zero to 12, and I picked 10 as the number of months I wanted. And during that time I was able to find a position at Maryland that would allow me to continue in effect, and that was a position at ISR (Institute for Systems Research). In truth, it was a good thing because I used to go up there about once a week; in theory, at NSF you get one day for your research.

Yost: I've heard that doesn't really play out.

Landwehr: Yes. That's Sunday or something. But I did make an effort, particularly in the beginning, to actually physical go up there to College Park, and since it's local, you can do that. I had an office there and I spent a day a week there, or a day every two weeks, and I would, among other things, have lunch with faculty. And it did give me, I think, a much better feeling for all the pressures that are on the faculty and what was going on, so

it was helpful. But I didn't really start a research program there, or anything like that. It was a convenient home.

Yost: Did you ever teach while this was happening?

Landwehr: No, I didn't teach any courses. I did actually teach at Maryland before all this happened. Cathy Meadows and John McLean and I taught a course one semester, back when we were all at NRL. But teaching, I couldn't even imagine being a program officer at NSF and teaching class, [laughs] I wouldn't do it. [After teaching at Purdue, I did teach an undergraduate CS course once or twice at Georgetown and a graduate level network security course for Virginia Tech, with David Goldschlag, so I know very well the effort it takes].

Yost: And you were chief program leader and program manager of ARDA, DTO, IARPA, from 2005 to 2009. Can you compare and contrast that program, and that role, with what you'd been doing at NSF?

Landwehr: Yes. So ARDA was Advanced Research and Development Activity. It was sort of an attempt in the intelligence community to start something that was DARPA-like and it got started in the early 2000s, I think. Some people saw it as just stealing some money off the NSA's budget or something like that. I don't think that was actually true, but there was controversy in the intelligence community about it. When my term at NSF was ending, my first term — typically, NSF won't let you be a rotator for more than three

years, sometimes four — and so I was there for four years, and by that time, I felt like I was ready for a change. I felt like I'd done some good things but I was, I think, worn out from that particular job. So when I went to ARDA, it turned out they were just getting a new director for the agency, which was Mike Macedonia, and since I was coming in at the same time as he was coming in, I thought it would be wise for me to talk to him before I got there. So I did call him up a little bit in advance and talk to him a little bit to see how he planned to run things. It sounded good to me and he seemed to have confidence in what I could do for the organization. So basically, I went up there — initially, the clearances had to come through and stuff, so I couldn't really get on their systems, but I could sit there with my own computer and draft stuff. So I used that initial couple of months to basically start on a program description of something I thought I'd like to do. And what's more, I had discovered that I actually had a couple of people who were there to help me [John Farrell and Lee Beausoleil], who were quite capable people, and so I thought I had died and gone to heaven. After NSF, to actually have people that were — NSF has good staff support and I wouldn't say anything bad against it — but it's not the same kind of staff support that you get at a DARPA or a IARPA, where the staff support is actually technical people, as well. So I benefitted greatly from that. I inherited a program that was there, so I had something to run, but I also had a budget, basically, and I had pretty good authority over that budget. So I was able to talk with people in the community [Fred Schneider and others] about what seemed like an area that would be helpful, or areas that would be helpful in how to frame a solicitation and got that started again, from ground zero. Well, it wasn't ground zero because there was a program there, but the program that was there was pursuing things that I didn't think were that

interesting, from my perspective, or that forward looking (intrusion detection and network forensics). Not unimportant things, but I thought I wanted to pursue information flow, and some other things in that line. So we put out a solicitation; we got a bunch of proposals in; made some selections; and that's how the program got started. We had a focus on accountable information flow, pursuing several different approaches, and also on managing large scale system configuration. That was a good time. [Subsequently, after helping to organize a series of workshops on privacy issues in the intelligence community, I developed a program to push ideas from private information retrieval to generate some practical results, called Automatic Privacy Protection, APP. Just before I left IARPA, I initiated a second program STONESOUP (Securely Taking On New Executable Software Of Uncertain Provenance – a writer for the defense press gave me the award for worst acronym ever for that, but I liked it). The idea was to provide some tools that an end user could use to tell if a piece of software (source or binary) was free of certain classes of vulnerabilities. My successor at IARPA, Konrad Vesey, very ably got STONESOUP off the ground and extended APP into a program called SPAR – Security and Privacy Assurance Research.]

[BREAK IN INTERVIEW]

Yost: So in 2009, you returned to NSF to run the trustworthy computing program for, was it three years?

Landwehr: Well, two years is what I stayed. I mean they would've let me stay longer.

Yost: Compared to the first time you were at NSF, had the budget expanded significantly?

Landwehr: Yes. I don't remember the numbers, but it had continued to go up. I'm going to forget the name, but it was the big; spending when, you know, we had the financial crisis in 2008 and then there was this act to [pause]

Yost: The stimulus.

Landwehr: Yes. The stimulus, right. I knew I'd forgotten the right word for that [it was ARRA, the American Recovery and Reinvestment Act of 2009], but that brought a huge wad of money specifically at NSF and some of that went into the program. So there was a lump sum addition and that had, I think, created some problems because they got a little behind because, really, a whole lot of extra work was dumped on them. I don't blame them for it, but when I came back, there was sort of remnants of that were still in the system so we had to get them cleaned out.

Yost: What did you see as NSF's priorities with regard to what you wanted to fund?

Landwehr: I guess, you know, different people describe NSF in different ways. I think I really do believe in the original model; in the idea that you fund what comes in that looks good; that the panels you recruit think are the most promising things to fund. I don't think

I would say that I had a strong agenda that I want to pick this or that, but rather I do have a bias, you know, in encouraging people to submit things that will bring some interesting results and maybe have a chance in the real world. Of course, NSF always has its two criteria: intellectual merit and broader impacts. The programs that I was involved in, cyber trust, and then subsequently it broadened out to include SBE and mathematics, and I think engineering, too, in Secure and Trustworthy Cyberspace. It always seemed to me that the whole point, in some sense, of these programs is to improve the trustworthiness of the cyberspace we live with, and that has a broad impact. So if you had a problem making a broad impact argument in your proposal, you were probably in the wrong program. Usually, of course, these arguments having to do with outreach and diversity, and things like that — which are all laudable objectives — but I think they're really there, originally, because there was a risk seen that NSF was just going to fund these ivory tower academics to do whatever they were interested in and there wouldn't be any benefits. But that never seemed to me like it should be an issue ever with any of the programs that I was involved with, and I don't think it was. I wanted to fund things that made a difference, you know, I'm in favor of those things. Of course, you get to recruit the panels, but usually you recruit the people whose opinions you liked, and they give you their opinions and you ought to listen to them. And so generally speaking, I did try to listen to them. Sometimes you think "I really wish this proposal had done better" but I would usually respect their opinions. I will say, early on, when we started funding large proposals in cyber trust, you know, one year we got a certain number of proposals that focused on particular areas. And if we thought that proposal didn't make it this year but that would be a good area, we would try to encourage people to resubmit and try to help



them make a stronger proposal. Some of those had to deal with national infrastructure, and voting, and things like that that we thought were important topic areas.

Yost: Are there certain projects that you funded that stand out that you feel really did achieve that role of major broader impacts, and if so, what are some of those?

Landwehr: This'll seem odd; I think there are many that I funded that had those kinds of impacts. They often happen a long time after you fund them and so it's a little hard sometimes to recall them. The two that I'm going to mention actually I funded when I was at IARPA. One of them, at least, came into NSF and when I arrived at what was then DTO, I had a certain amount of money that had been left over, sort of, from the previous year and so I had some money with which I could fund a few things. So I asked my friends at NSF if they had any things that they were unable to fund that they might be interested in having me fund. I got several proposals, and I looked them over, and my staff looked them over, and we picked a couple. One of them was a proposal from Nick McKeown that ended up in the open flow architecture, which has made a big difference and is, I hope, going to continue. There was also some work that was applying model checking to configuration control. That one actually came out of Telcordia, at the time, and I think that one's made some difference. There were others, you know, work on information flow; there's all kinds of them that made good contributions, so I sort of hate to single out particular ones. Those two strike me because those were really discretionary funds that I had, initially; I mean, I made a choice about them.

Yost: Beginning before, but overlapping your time at NSF, you became Editor-In-Chief of *IEEE Security and Privacy*. Can you talk about your vision for the journal when you took that on, and what you thought had achieved before, and what you wanted to do?

Landwehr: [Laughs.] I think my attitude is not so much having a vision; actually, I guess I applied for the position before I got to DTO, but I didn't actually assume the position until after I was there. So I was an associate editor when I was at NSF, but I became editor-in-chief after I was at DTO. I did it partly because I thought it would be good for me to have a presence in the community that way. I thought that George had gotten things started, and my real ambition was to keep things going. [Laughs.] A lot of the work for magazines like that is really just to be sure the thing keeps going, and to be sure that you've got people around you that can generate the columns, the content, that you want to have. I won't say that I had a specific vision as to some place I wanted to take it. I wanted to make it interesting and grow the readership, and that's challenging enough.

Yost: Before we wrap up, I have a handful of broader questions. What are some of the most important things that you think computer security researchers today haven't learned from the past and could benefit from?

Landwehr: [Pause] That's a tough question. I think that, to some degree, you seem to see ideas such as we see in virtual machines and so on, reappearing, but I don't know if that's because people weren't aware of them or, given the way that computing architecture has evolved, they've become relevant again in new ways. The hard thing seems to be to make

the connections with industry and maybe among members of the community. I'd say the broadening of the scope of the research program at NSF, for example, is good in that I think the problem is broad. It involves people, it involves economics, and so we need to deal with those things but I don't see the community speaking out very strongly with one voice as to how we ought to be building systems to make them less vulnerable. I think that's hard because instead, every academic has to grow his own reputation, and with his own ideas, and so on; and I think one of the things that I've tried to precipitate now is some unified agreement at some low level as to what's the proper building code for infrastructure software systems so that we don't have things happen such as happened within the last two weeks with Heartbleed, where you have a piece of infrastructure that everyone — most of the public, anyway — is completely unaware of and certainly doesn't think about, and it turns out it has a fairly elementary mistake in it, but made a whole lot of things vulnerable and nobody noticed the error. It just seems to me that we shouldn't be in that kind of state and the technical community needs to act together to deal with that.

Yost: In 2005, in a panel discussion — I believe this was in *IEEE Security and Privacy* but I'm not certain — you talked about the rediscovery of virtual machines as an isolation mechanism. Can you tell me a bit more about that rediscovery, were there are some principal people that were advocating returning to an older idea?

Landwehr: Oh, boy, [laughs] I've lost the context completely from whatever panel that was. It's obviously all over the place in cloud computing. I don't know where it was in

2005, specifically; I do remember funding, early on, a proposal or two at NSF that involved, I think it was at Michigan, some work on what's called on subvert and revert, which was basically putting a layer underneath the operating system to monitor it, but I'm afraid I'm not giving you a very good answer on that one.

Yost: Looking back now at the early on history of computer security, what do you see as the various accomplishments and what areas do you see as some particular challenges or areas where the research community has kind of fallen short?

Landwehr: Well, this will seem a little odd. There's no question that there are great accomplishments in cryptography. Public Key-cryptography is a major innovation, and in the last few years, the results on homomorphic encryption, which are not yet in the practical realm but might be coming that way, I think are interesting; and just in the past year, in the area of obfuscation. So I think there's very interesting developments that have gone on, on the cryptographic side. The notion of security kernels, I think, the idea that we now have proofs, that you can enforce the same policies with a distributed kernel as with a centralized one, I think that's an advance. I think that probably continues to be a more likely way that we're going to get the security that we want. I think that there's been advances in the ability to monitor information flows that give me hope that the kinds of policies that Dorothy Denning's original work pointed to, and which to me, continued to be the more intuitive way to look at systems in the future, provides some hope that we might actually have systems where we can describe the flows of information

to people and to force those flows in ways that people agreed are what they want. Those are very specific technical things but I think that's how I might describe it.

Yost: And areas that haven't progressed as well as you or others might've hoped?

Landwehr: The fact that we still seem to be at the mercy of relatively low level implementation flaws is disappointing, that we're still using languages and tools that don't rule those things out, and the arguments for not using them haven't been engaged forcefully enough. There's a talk that I've been giving on we need a "building code" for building code. One of the examples I use at the end of that is actually a review in the *Wall Street Journal* of a hybrid car, a Lexus, a fairly sizeable hybrid car and the reviewer says basically that you know, a few years ago, people were telling us that these standards for fuel efficiency were gonna lead us to a lot of cars that were very unsafe and not fun to drive. In fact, this car is safer and more fuel efficient, and more fun to drive than anything that was around here 10 years ago. And so that just wasn't that; those bad futures didn't happen and in fact, I would argue that it's the constraints on the process that can make engineers think harder about the problem and lead us to a better future. And I think we need to be a little more willing to impose some constraints. How to get people to do that, you know, how to get them to say okay, we're going to give up 10 percent in order to be sure that we don't have a Heartbleed in the next year, because it's a social question, in some cases, and we have to deal with that.

Yost: Are there areas I haven't asked questions that you'd like to discuss? Or areas we have touched upon that you'd like to say more?

Landwehr: I think it's just more like leveraging the last question, I guess, it's another area I've been in is trying to come up with creative types of competition. I think that right now, there's a lot of energy, and probably there's been a lot of progress in the area of intrusion detection and patching, certainly, and there is a lot of public attention on a lot of competitions in which students are pitted against each other and trying to break into systems and defend them. I'm interested in competitions that would teach us how to build things better. And in particular, I think right now, the tools that we use to build these systems, in general, make it too hard for ordinary people to build systems that lack vulnerabilities. That was a contorted sentence, but what I mean is, it's as though we have knives but they don't have any handles and so we cut ourselves a lot on the tools that we're using. I think we need to stimulate the development of tools that ordinary people can use and result in systems that are not perfectly secure, perhaps, but are not full of vulnerabilities, the way many of our systems seem to be today, even when people try to take good care, and I think we need to learn how to do that. So one of the competitions I've been working with people on is a competition that would result in the development of, or the identification of chains of tools that ordinary people can use to develop systems with good security properties. I think we also need to develop ways to explain technology to the public that uses it, but doesn't really understand it; I think that's important, too.

Yost: Okay, well thank you so much.

Landwehr: It's been my pleasure.