

An Interview with  
GEORGE STRAWN

OH 475

Conducted by Jeffrey R. Yost

on

10 March 2010

Arlington, Virginia

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

## George Strawn Interview

10 March 2010

Oral History 475

### Abstract

With support from the National Science Foundation (Grant No. 0811988, “Designing and Using FastLane: Distilling Lessons for Cyberinfrastructures”) CBI researchers Jeffrey Yost and Thomas Misa conducted oral history interviews with 70 NSF staff members as well as numerous additional interviews during 29 university site visits. An overview of the project is available at <[www.cbi.umn.edu/oh/fastlane/](http://www.cbi.umn.edu/oh/fastlane/)> and a complete set of 643 publicly available interviews is at <[dx.doi.org/10.13020/D6RG6B](https://dx.doi.org/10.13020/D6RG6B)>. Here on the CBI oral history database is a selection of notable NSF staff including Joseph F. Burt, Jean Feldman, C. Suzanne Iacono, Constance McLindon, Carolyn L. Miller, Paul Morris, Andrea T. Norris, Erika Rissi, Craig Robinson, Mary F. Santonastasso, Rich Schneider, Frank P. Scioli, Beverly Sherman, George Strawn, and Frederic J. Wendling. Topics common to many of the interviews include the design and development of the NSF’s FastLane computer system, interactions with users, e-government initiatives, grants management practices, peer review, and NSF policies and practices. These interviews span a wide range of NSF staff, from program officers to senior managers.

George Strawn was NSF’s first Chief Information Officer and a keen observer of interagency activities, including the Grants Management Line of Business that led to Research.gov.

Yost: My name is Jeffrey Yost, from the University of Minnesota, and I'm here this afternoon with George Strawn at the National Science Foundation. It's March 10, 2010. Dr. Strawn, can you begin by telling me when you came to NSF, which directorate, and briefly describe the different positions you've held at the National Science Foundation over the years?

Strawn: Surely. I first came to NSF in 1991 to serve a term as the NSFNet Program Officer. NSF was building the thing that we called NSFNet, then, which used the internet protocols, which ended up taking the DARPA Arpanet protocols and turning a small experiment into a major infrastructure for higher ed, which then caused the private sector to look at it and say there's money to be made here. Therefore, we were the bridge between the Arpanet experiment and the commercialized Internet, which emerged in the mid-'90s. So I helped in that privatization and commercialization phase in the '91-'95 timeframe, as a Program Officer. And I actually came back to NSF again, having just left, as the Division Director for Networking in the CISE Directorate in '95. I held that position for about three years and then I became the Deputy Assistant Director for CISE [NSF directorate for Computer and Information Science and Engineering]; held that position for about three years; then I was Acting Assistant Director of CISE for a year between Roug Nabyji's [?] reign and Peter Freeman's. And when he came I became the NSF CIO and served a term of about seven years as NSF CIO; and then last fall went on detail from NSF to my current position of the NITRD NCO, Networking and Information Technology Research and Development, which is an interagency program run by the Office of Science and Technology Policy. NSF hosts the national coordinating office for

the NITRD program, so I'm Director of the NCO and Co-chair of the NITRD subcommittee. So that's a long history for somebody who thought he was taking a temporary two-year sabbatical from Iowa State and expected to go back by '93 and be done with government service.

Yost: Were you involved in the EXPRES project or were you aware of that project as it was going on?

Strawn: I was quite aware of it because it had just been shut down, I think, or concluded let's say, in the '91 timeframe. Two or three of the members of the Networking Division, which I joined as Program Officer, had recently come from the EXPRES project and so I heard many private opinions expressed about the various issues that were surrounding the EXPRES project and what was going on, and of course, they were somewhat sorry to see it concluded because they had been committed to its success. I think I sort of summarize all their feelings and what I know about the thing is, it was a wonderful experiment but it was done with pre-internet technologies so the tools just weren't right. I mean, EDI, Electronic Data Interchange, was probably not the right tool for academics to utilize in a proposal mechanism. But my own opinion was we probably learned enough through the EXPRES project that that helped the success of the FastLane project four years later. So it was a good example of research where if you don't fail some of the time, you're not trying hard enough.

Yost: Would you mind telling me the names of those two individuals?

Strawn: Don Mitchell was one; maybe Dave Stout; maybe Doug van Bellingham. I think at least two out of those three were involved and I guess you know the name [James Morris? Dan Atkins?] of the Principal Investigator of the EXPRES program? A guy from the Applied Math Division; I met him once or twice. I can't even say his name right now; you could probably tell me, but I can't. So those were my connections.

Yost: Okay. Was there a relationship between your team at NSFNet and the early planning for FastLane?

Strawn: Not that I was aware of; on the other hand, let me tell you what I viewed was the connection. We had a Division Director in the CISE Directorate for Advanced Scientific Computing. I'll have senior moments all over the place here as I try to remember names. This guy just retired from NSF a year ago. You may know his name. He retired as the Division Director of the Metallurgical Sciences Program in the MPS Directorate. I'll probably say his name in a minute if you can't prompt me. But during the reign of CISE back in the '93-'94 timeframe, he left that Division Director position in CISE and for one year, maybe at most two years, became the Division Director of Information Systems in NSF. While he was there, he wrote a white paper, which proposed that NSF should take another crack at the electronic proposal receipt problem. He gave me a copy of that paper a few years ago and I read it, and said yes sir, this is certainly where that project began. Then he moved along and Fred Wendling took the job. About the time Fred took the DIS

job, the Mosaic browser was invented at one of our supercomputing centers. Sorry that it competed with Gopher from the University of Minnesota.

Yost: It was far more successful; it was graphic based.

Strawn: Well, I'll insert a little war story here. I was Director of the Iowa State Computation Center at that time, sort of splitting my time between NSF and there. Before Mosaic, I called in some of my technical advisors and said do you think the World Wide Web will replace Gopher as the information mechanism of choice? And to a person, they said no, it'll never happen, because it's too complicated. Gopher has the right level of simplicity that will endear itself to the public and the World Wide Web is too complicated. Well, practically while they were talking, two undergraduates at Illinois were designing a browser that hid the complexity. So once the complexity of html and all that stuff was properly hidden, the mere fact that you could store information at the branching nodes of the tree, rather than just at the root nodes of the tree, as I recall was the case in Gopher, the rest is history, of course. It turned out to be the right complexity with the majority of the complexity hidden from view. Well, Fred was a smart guy and if you haven't interviewed him (pause)

Yost: We have a very extensive — I think about five-hour — interview with him.

Strawn: Good, because I think he's one of the major figures and I've interviewed him too, so I'll tell you what he told me and you can see how that compares with what he told you.

He went out to Illinois to talk with the people and say, well, do you think we could use the browser and the internet technology as the basis of an electronic proposal receipt capability? And the advice he said he received was oh, no, oh no, it's too early, too early; couldn't do that. Well, Fred was not to be discouraged; went back and proceeded to fire up the project in the sense that he first wrote a detailed white paper, and the previous white paper that I alluded to, was conceptual. And now Fred's white paper, which I've also seen a copy of, and you may have too; sort of specified getting down to brass tacks how the web would be used to do such a thing. I think Connie McLindon was his boss at the time, and she's also a clear-thinking agent. Connie and Fred, as the report I had it, went to Neal Lane, who was Director at the time, in effect said we'd like to take another crack at this problem, what do you think? And, according to Fred's report back to me, Neal showed his understanding of the federal bureaucracy and what dysfunctions it's sometimes prey to when he gave them two answers; three answers. He said, yes, I want you to try it. I'll give you more money to work on it. And it's alright if you fail.

Now, the worst thing you can do in the federal government is fail, right? That's why we're late adopters of technologies almost all the time. So here was permission to take a risk and do something with a pre-statement that I know this an avant garde project, it's alright if you fail. Not that he wanted them to fail, but he gave them permission to go ahead and try something risky. Beautiful example of the government working the way it should, because of smart worker bees and an enlightened king bee.

This was all happening about the time that I came back as Division Director so in either late '95 or early '96 a call came out from DIS saying oh, we're going to pilot a new electronic proposal device receiving proposals over the internet. Well, I said to our division people, hey, that's our network and our division, why shouldn't we be the first division in the Foundation to pilot this electronic process. They agreed with me and so we put our name in the hat. So if we weren't the first we were certainly a positive responder to the first round of guinea pig divisions who would try the new FastLane. So we did and, of course, as I'm sure you know, by the late '90s we were in full progress. If I'm remembering right, I'm thinking by '99 we were 95 percent electronic proposals, or something like that. So, extremely productive electronic [pause]

Yost: Prior to the time it became mandatory throughout NSF in October 2000, it was mandatory in a number of Divisions within CISE.

Strawn: Yup, and like I say, we were certainly right up front there, in the Networking Division, and CISE as a whole would certainly look at this as an application of their domain, which they should support and did support. So the late '90s was a very productive period; based on the intellectual ferment of the early '90s; based on the previous attempt with the EXPRES project in the late '80s.

Yost: Can you talk about the exchange of information and the role that CISE played in an advisory capacity, it being a pilot with this new technology and testing this new technology?



Strawn: Well, next you're going to ask me if the Department of Architecture at a university advises the physical plant on planning for new buildings. Not likely; or not extensively; and I would say the same was here true. Now, about that time, shortly after I became Division Director, the then-CIO fired up a CIO advisory group and I was a charter member of that group, and remained a member of that group until I became a CIO and became a continuing member on the other side of the table. Especially in the early days, it was a very high level review of things at that formal advisory group. And at the program officer level, the Division of Information Services, DIS, has a pretty good record, in my opinion, of putting together advisory user groups, assembling program officers, or program assistants, or other operatives who have to use these systems; getting together once a month and letting them gripe about the things that aren't working well and compliment them on things that are working well, and then sort of, what are the highest priorities for making new adjustment, so on and so forth. My recollection was that CISE didn't play a special role in that; that program officers and others from around the Foundation served as user advisors to DIS as they proceeded with the development.

Yost: Do you recall what types of comments you were hearing from the research community when FastLane was launched?

Strawn: I would say the absolutely normal comments when anybody is forced to do any change. I for years have had the opinion that all change — good, bad, or neutral — creates at least as much heat as light because people don't like the change. Then you hear,

oh, I have to learn something new and it's taking time away from the laboratory, and now I'm having to screw around to learn how to use this new system; and by the way, it doesn't work well when I do X, Y, and Z. But mostly, those were the sort of comments that I've heard with every type of change I've been a part of introducing in IT since 1960 and the level of comments told me things were going pretty well. By the late '90s, which isn't that long a time for the development of a major IT system, things were working well enough that they weren't griping. Then people were starting to say, gee, isn't it a shame that other agencies don't do this same thing. As you well know, by the early aughts [2000s], NSF was saying all right, now that we're receiving all the proposals electronically, we'll develop a complementary system, which we ended up called E-jacket, or electronic jacket, which sent us down the road to do the internal processing of turning grants into either proposals or declinations in a much more automated fashion, pushing bits around the Foundation rather than pushing paper around the Foundation. Now that's pretty well complete. It was held up for a few years by inadequate funding but fortunately, we've been able to change the funding model, which should keep that type of holdup from occurring again in the future on future products.

Yost: I understand that you were co-chair of an interagency federal networking council from the mid-'90s to '97.

Strawn: Which is sort of the antecedent of this group I'm in right now, by the way, the Federal Networking Council morphed into the large networking subgroup of this NITRD

group. And then I became co-chair of the LSN group after having chaired the Federal Networking Council. So yes, you got that history right.

Yost: Was there any kind of discussion within this group of the different work that was being done on grant submission systems using the web?

Strawn: I don't recall any. Both the Federal Networking Council and the LSN after it, tended to focus on scientific use of the network as opposed to scientific administrative use of the network. So, if there was any, it was incidental enough that I don't remember it.

Yost: What was the impact of FastLane on NSF work processes?

Strawn: May I talk about FastLane plus E-jacket, as the whole electronic milieu?

Yost: Yes.

Strawn: I think roughly the same except at a more gargantuan level, as e-mail put secretaries out of business and made us write our own memos. Well, we got them written faster, and they went out quicker, and so on and so forth; and we had to learn new etiquette, and whatever, but the number of administrative support staff to do the communication process went down as we started doing more ourselves. And that probably made us a little busier because we were taking on a job that we had offloaded to support staff before then. I think, in some regard that was a result of automating the

proposal process, too. Things could move faster. Automation always gives you the chance of doing stuff faster than you did before. Well that means you have to work harder to keep up with the faster that the new machines are working; think of Charlie Chaplin's "Modern Times," movie where he's trying to keep up with the assembly line. It certainly changed the balance of work between what program assistants and administrative folks did versus what the program officer did. The program officer started doing more things, in many cases, for themselves. We probably stretched the support personnel more to learn the new automated procedures, whereas the professional staff, the program officers, you could sort of hold them accountable for understanding the high tech approach more quickly. I think that ultimately settled down; I just thing there was a bigger learning curve for the program support staff. And, of course, more gripes, when the program support staff would sort of take it as a given and learn how to use it; and the program officer would say that's a stupid way to do that, why don't you do X, Y, and Z instead? Then you get good suggestions for improvements, even if they weren't always expressed in the politest of terms.

One thing that the Foundation had been criticized for in previous times was the length of time it took from submission of a proposal until the proposer knew whether it was going to be funded or not. That was in good measure because there were so many manual steps involved. So, by automating away many of those manual steps and by changing some business processes, the Foundation was able to establish a goal, and achieve the goal of processing more than 90 percent of their proposals in six months, if I remember the figures correctly. And we did accomplish that. It's not clear to me that having FastLane

and then E-jacket were required in order for us to achieve that type of a goal. Now, there are all sorts of unexpected consequences. Now an unsuccessful proposal would be resubmitted to us twice a year rather than just once a year. So in one sense, it might have made a little more work because you got it back faster and made a few changes, put it back in, so on and so forth. But that's just one of the unanticipated consequences that you have to anticipate; that there will be unusual things like that.

I don't think we know yet all of the important effects that will ultimately arise because of FastLane and E-jacket. The first step of any automation project is just start doing electronically what you used to do manually. And then after you get it done, you say, oh, wait, there are new possibilities of things that we could do that we couldn't even think about in the old days. Well now we've started doing that with this second generation FastLane E-jacket system called Research.gov. You should have some talks with people about that if you want to look forward as well as looking back. Research.gov is introducing capabilities that [we] couldn't have been dreamed about in the paper days and therefore, weren't thought about in the first generation of electronic systems. For example, we now have an electronic database of all proposals ever submitted to the Foundation in the last 20 years, or something like that; a very long time. Now program officers can search that database; so it's like a Google search within our specific domain of proposals. That's pretty good and in Research.gov that was first announced as just a subfeature of another "Find Reviewer" project. Finally, it was clear that it was so important to the program officers it should be a standalone facility and it's unbundled from the "Find Reviewer." What's "Find Reviewer?" Finding reviewers is more

complicated now that we have more proposals and the communities are bigger, and so the program officers may not know everybody in the community who could be a reviewer.

So when a proposal comes in, we've got to look at that proposal, and one source of reviewers is to find people who have submitted similar proposals. So if you can measure the similarity of that proposal to previously submitted proposals and get back the name of those investigators, they can be candidates to be panelists to help review proposals such as this. That's a wonderful electronic improvement over just expecting the program officers to know everybody who might be appropriate within their community of interest.

So, we've already done some experimentation with using electronic methods to take electronic proposals and group them together. I mean, one of the big jobs that a program officer has when, say, you get a hundred or in huge cases, a thousand proposals, you have to break these down into groups of 40 or 50 proposals for independent panels. How do you make sure you put coherent subgroups together so that you have the best likelihood of reasonable reviews? Well, there are data mining techniques that allow you to analyze those proposals in such a way that you can get some pretty good suggestions or choices for how the panels can be put together. So clustering proposals for a given panel, finding reviewers for the various panels, we're just at the front end. We've done extensive experimentation and we're at the front end of being able to add those services to the services which are automated and therefore, don't just require manual knowledge of program officers. I'm convinced that there will be even more important dimensions down the road as we master the art of data mining, in the sense of textual data, not tabular data, where we analyze the proposal submitted to understand what the shape of the interests are

of the research community at a given time, and follow it longitudinally over time. NSF did give an award to a researcher, at Arizona, I think, to study the 10-year history, at that time, of NSF support for nanoscience and technology. And that had to be done with some creativity, because at the beginning of that period, the phrases "nanoscience" and "nanotechnology" weren't even in existence. But yet, by proper text mining, you could figure it out. You could also, by proposal analysis, you could sort of define subareas that are different from what a human might look at and then you could follow those subareas shrinking and growing and shrinking in interest, as the community changed directions in looking at things. Over time; this is going to give us the ability to sharpen up the way we do our solicitations for proposals; going to be able to sharpen up our description to the decision and policymakers downtown of what's going on in the research community, and where the community is working hardest, and what we think the opportunities are, and so on. So I think the uses of electronic proposal data and then, subsequently, in concert with the electronic research articles that result from those proposals, which, as you know, is a big brouhaha in society right now about open access, public access to scientific literature. That will happen; we will have more public access. We then be able to compare and contrast the proposals with what the investigators proposed to do; what did they end up doing. Science takes a lot of unexpected turns. It wouldn't be science if you knew what you were doing, or knew the end when you started out. So, we even have a program at NSF called The Science of Science Policy, which is attempting to make this whole business of how things work out more of a social science, let's say. All this is almost impossible without electronic proposal data and subsequently, electronic research article data. Extremely important.

Yost: And are the data mining tools currently available to program officers?

Strawn: Yes, this Research.gov with its various subservices, the services and tools that are available, has been coming into existence over the last year or two. You could speak with Andrea Norris and her people in the Division of Information Systems for a forward look at what second generation FastLane looks like. That'd be a nice adjunct to your paper, if you were to look forward at this point, then you guys can submit another proposal to us in 10 years, come back and see if it turned out the way we thought it was going to turn out.

Yost: And is this largely making PARS obsolete?

Strawn: Yes. PARS will go, as part of the new Research.gov. By the way, Research.gov was part of the previous administration's e-gov projects, where one of the things they wanted to do was to get the agencies to work together. One of the first and easiest to implement was electronic payroll, where they said we should not have every agency in the federal government doing its own payroll. So they eventually chose four agencies and said to all the other agencies, now you choose to have payroll done at one of those four. We chose Department of Interior and their Denver processing center, and for the past five years my paycheck has come from Denver, rather than from here.



One of the other initial projects was, well, we should have a FastLane for everybody. That got called Grants.gov. Grants.gov was an e-government project for all of the other agencies to submit electronic proposal information. It has suffered from some implementation problems. First place, in my opinion, was the overambitious problem of grouping together competitive research grants and block grants to states. Block grants to states are about 10 times the money that competitive research grants are so heavy focus was put on what was required to give block grants to states, and therefore, they didn't really find that our comments to them that Macintosh support was necessary, to be pertinent because they didn't know of any state government where Macintoshes were in high use. We told them that if you try to do a competitive research grant to university people and Macintosh support wasn't available, they'd hear about it. Well, NIH eventually, using that system — no Macintosh support — they heard about it at high decibel levels. And then alright, the next version of Grants.gov will have Macintosh support. And so, what about UNIX? They said well, what's UNIX? And so it continues to have issues like that. Grants.gov has some support in the academic community and the idea is great, especially at large universities that have sponsored research offices like Minnesota, where you've got 17 different formats and that you have to get on top of all those formats and keep your budgets right; if the federal government could ever talk with one language on that subject, you would greatly facilitate that research universities' business. Rather than using a browser approach, which is the way that FastLane had worked, Grants.gov uses a batch approach. NSF has used it in the sense that if a proposal is written according to Grants.gov specification, it's sent to the Grants.gov office, which then forwards it to NSF. Then we slip it into our system as if it were written by a

FastLane process, but since you haven't had any interactivity part you find all sorts of missing commas and this and that and the other, and it always gets bounced. And then since they submitted it the last day, they missed the deadline and there's all sorts of terrible griping about that. But MIT and maybe a couple of other universities are working on a local module that could go on campus and do all the preprocessing the way Grants.gov does interactively, and so some of the universities are hopeful that that still will come about one day, so that they can preprocess proposals locally, and when they send it off to Grants.gov it will be pretty likely to pass muster.

Yost: Was there ever a possibility that FastLane would become the model for different agencies and if so, why did that not occur?

Strawn: That's a good lead into the rest of my story on the Grants.gov thing. I wasn't directly involved with those discussions at that time, but I'll give you my impression of what probably happened. We probably looked and said holy smokes, we designed FastLane to fit in the size and scale of our proposals, if we all of a sudden have 10 times more proposals because all of the other agencies are using FastLane too, we might sink the thing. We probably didn't write it to scale up to be much greater than what we're already doing because, of course, as our own use grew we had to keep modifying it and tuning it and getting rid of bottlenecks and so forth. So I think we were probably scared to do it, just for fear of being not able to perform well. And there were the normal feelings of, well, that's not invented here and we've got a new office and we can do it over and we can do it right, and so on and so forth. Only trouble was, in my opinion at

least, one of the troubles was that since it was an administration project it had to be done on a political time scale rather than a technical time scale, which usually means it's never time to do it right and we'll always take time to do it over when later times come along. So that helped plague its issues.

When we began the Research.gov second generation stuff, two things happened. First of all, we were asked by the administration to take the lead of a Grants Management Line of Business, e-grant project. That is to do something like our E-jacket but make it available for multiple agency participation. So we concocted the very true and believable story, well alright, E-jacket and FastLane and are getting old, we need a second generation of that technology and as we do the second generation we'll make it available for multiple agency participation. Now, that sold the fact that multiple agency participation was possible, gave us approval from OMB and OSTP to go ahead; the fact it was going to be a second generation of the technology already serving NSF made the NSF leadership comfortable with a second generation of those services. This, by the way, would be with a prospect of fee for service, so the agencies that would use [it] would pay their share of the use and NSF wouldn't be burdened with extra expenses. So those were the circumstance under which Research.gov began several years ago, with the admonition, don't duplicate the Grants.gov business. So the two areas called "Find and Apply," find funding opportunities and apply for a funding opportunity, they said don't duplicate those because everybody should use Grants.gov. Well then, maybe just a year ago, the new administration said, well, we're not sure how soon Grants.gov is really going to be competitive in the research grants area so why don't you go ahead and add a "Find and

Apply" capability in your second generation system as well, and then people can choose to use either that or if Grants.gov has come along, they can use that. So in fact, very belatedly, we will probably offer "Find and Apply" services in Research.gov, similar to what exists in FastLane and what exists in Grants.gov.

It's a long project; big project. Since we're just started, a few years into Research.gov, it'll be mid-decade before we really see the shape and scope, and what's working, and what we have to change, and so on and so forth. But we are, I think, well begun on the second generation of automation.

Yost: How did FastLane and, I guess, more specifically E-jacket, impact, if at all, internal reporting within NSF?

Strawn: Well, I guess there's different types of internal reporting. Let's also talk about the fact that PIs are supposed to submit final reports, which has an external source but then becomes an internal document. Under old time manual procedures, some directorates were vigilant about that, others were not. And, in fact, when FastLane initially came up, it eventually provided a facility to receive the final report electronically. But just like in the manual days, it was left to Program Officer recognition to observe that the report had come in and was satisfactory. The key is the official policy that you can't get another grant from NSF until you've submitted the final project report from your previous grant. But there weren't any automated mechanisms until the IG [Inspector General] looked and said, humph, we see a lot of grants here without final project reports, so then DIS was

authorized to put a hold and make it impossible to give an investigator a new award until the final project report had been submitted for previous awards. That caused a little angst from time to time because now all of a sudden they had to do in reality what they had had to do in theory before. We're now well over that; now people know darn well that they have to submit their final reports.

Another issue arose. Well, those final reports, shouldn't that be public information, the IG said, and others echoed. Why don't you make that public? Well, that was understood to be a private communication between the PI and the PD saying what went well, what went poorly; it's not peer-reviewed, it's not part of the peer reviewed science, it's a communication between the two. In this era of transparency and open government the final report has been redefined in a multi-agency capacity, which will have a public part and a private part. There will be a possibility of keeping some more confidential and some less confidential, and the less confidential part will go on a public website. We've already started that. We have an award search database where, initially we put up all the particulars of an awarded proposal; who was the PI; how much is the award; what's the award abstract. That turned out to be an extremely popular website, in terms of people looking in and wanting to see what NSF is funding. Then people said well, how do you know any scientific articles are being published from these given awards? Why don't you make available the titles and abstracts of the articles that resulted from those awards? We went through another bit of huffing and p-; again, this is something that could never have been done — could never have been done *usefully* — in manual days. So now on our award search database you will find an average of 5.1 articles referred to, with each

funded project, and we've done our transparency and given people a result to look in that fashion. Of course, since the scientific literature isn't still open, you get to see the title and maybe the abstract, and you'll have to pay for the article if you want to see more; or go to a library that already has a subscription. But that's the type of reporting and so on and so forth that's sort of internal and external that has been facilitated by electronics. Just the mere fact that we have the proposal abstracts electronically because they're part of the FastLane input, now we can put up that type of a database.

Yost: Did it also have an impact on how managers, in terms of transparency for managers, with program officers, and work flow?

Strawn: Probably. But I was out of that business by that time and pushing other stacks of paper around or doing other types of things so I don't have enough personal experience that I can give you other than, a yeah, that's probable.

Yost: Do you think it changed the way Program Officers interacted with their Program Officer peers? Specifically E-jacket.

Strawn: Insofar as the electronic thing sort of changes the way everybody acts and reacts and does more and so forth, one of the big issues we've always had is, what about interdisciplinary calls for proposals; or what about so-called orphan proposals that you're not sure whether it goes in this Directorate or that Division and you're trying to decide who should take responsibility for it. The analysis, the data mining of proposals that I

talked about earlier, was used to great effect in some big interdisciplinary projects where a proposal had multiple disciplines represented. ITR, the Information Technology Research project, which had to propose making it a research advance in some area, using IT, but a research advance in the IT that was used to support the other one. So, not just as a tool, but as tool development and subject development. The data mining, you might speak with Suzi Iacono.

Yost: I actually have an appointment to meet with her this trip.

Strawn: Suzi ran the ITR project, I think, the year they had more than a thousand proposals and she said without the data mining capabilities, which we would call the Pisani software, because Mike Pisani was the Division Director who had a graduate student develop it for us; she said she might have sunk without that, in terms of being able to group proposals into panels and to figure out what to do with orphan proposals — is it really more like this group or is it more like that group — and that was a great facilitator in terms of trying to help Program Officers work as a team both in interdisciplinary and subdisciplinary environments. Again, the true promise of that is yet to come but we've seen enough pilots that I'm willing to bet on it.

Yost: Do you feel that FastLane changed PIs interactions with NSF staff in any meaningful ways and if so, how?

Strawn: Well, the most meaningful way that I'm aware of is the rapid turnaround. We're, roughly speaking, able to get information back to a proposer in six months rather than 12 months. But, of course, our volume of proposals keeps going up year by year; I think we're up to maybe 50,000 or so proposals a year that we receive, with only 10,000 awards being made so it's a big job for the relatively small staff that we have. And the bigger the job gets and the busier the Program Officers become, the more formal the interactions have to be with the outside community; the more it's e-mail and not telephone; and short answers, not long answers. Of course, we'll always be concerned with not giving anybody an unfair advantage by giving them too much information, or that sort of thing. So my guess is faster interactions but sometimes suffering delays where a Program Officer is just too busy to get information back to a proposer.

Now, an awardee, well, there are different standards in different Directorates and with different sizes of awards. If you have a single investigator award in the average Directorate, once you've got your award you really don't have much need to interact with the Program Officer. You can even make some budgetary changes without having to get permission, so we were able to automate that. That takes the Program Officer out of the loop [more] than might have been in manual days. And for a small award by NSF standards, a quarter-million-dollar three-year award, let's say, the way we do an evaluation on that award is what does the panel think of the work that you did on that award when you come back to apply for another grant? You have to show what you've done; you have to show the articles that you published; and the results you obtained, and so on and so forth. If you can't show that you made good use of the previous money, your



chance of getting an award the second time around are hugely reduced. But if you have a large award — and NSF has been pushed into doing more larger awards, multi-million-dollar awards over the years, as they want us to take on multi-investigator and multi-disciplinary work as well as single investigator work — that requires a lot more award management and oversight than just waiting and seeing how it turns out and evaluate it when they apply for something else. It's not being a good steward of the public purse, if you take a \$15 million project and you just wait and see how it happens. So I think NSF is still trying to staff up with enough Program Officers, or enough more automated procedures, that post-award management can keep up with the larger awards that we're being asked to produce. My guess is that automation will continue to assist us to meet those responsibilities better.

Yost: What do you see as the greatest challenges to implementing FastLane at NSF?

Strawn: I'll make a general statement about implementing software in the governmental view, and I will say that NSF has met this challenge better than many hapless agencies have, is having to work through contractors. You can't hire any more government people; you have to farm out to contractors, the development. When I got here in the early '90s I think we had 60 people in IT, federal people in IT, now we have 50. But we also have 200-300 contractors. So all the growth, as IT has taken over the way we do business, all the growth is in contractors. So unless you have people that are smart enough to manage contractors and get work out of them and have them assume their responsibilities; they have many more lawyers than we do, limiting their own risk and liability. We do

defensive programming just like the medical field does defensive medicine, right? You make sure that you don't take on any risk that you weren't supposed to, and therefore, each one has a little compartmentalized section. Well we have the whole thing. So you just have to: a) have good contracting to get the right vehicles in place, and b) know how to put the contractors feet in the fire to implement stuff in a right way and not get fleeced in the process. I'm happy to say I think NSF's been pretty good at that. Not perfectly, but you hear about billion-dollar projects around the federal government that have to be shut down, standing the billion-dollar investment because they've gotten into a hopeless situation.

Yost: The major computer system of the CIA.

Strawn: CIA, IRS, FAA, you know, these are legion, and I think the reason is because software projects seem to be much harder to do things right than hardware projects. Because we know once something is built in hardware it can't be changed very easy so we're really careful about design specs and checking stuff out along the way. Software, everybody says, well, you know, we can change it. Well, yeah, you can change it, and change orders cost like hell. So, I think, most of these big, big projects they started off with half-understood ideas of what they wanted and then, oh, now we know what we want, and start submitting change orders, and pretty soon the change orders cost more than the original order did. Sometimes when I'm depressed I say software engineering is more an aspiration than a discipline, especially when you complicate matters by working through third parties under strict contractual obligations. So any government software is

complicated by the contractual arrangements of working through the private sector, in my opinion.

Yost: Are there rules on contracting that...

Strawn: Are there rules!

Yost: ...if post-problems with the FastLane project, I think it was in 2004, CompuWare's contract ended.

Strawn: Yes, we were moving; as we're doing bigger time applications, we're moving from smaller contractors to bigger contractors. Well, that has its pluses and its minuses; but that's more or less what we've done; and been more or less successful at it. The Federal Acquisition Regulations, the FARs; when I'm feeling mischievous like I did a year ago in a conference with vendors to the federal government and CIOs, I said I know how to ruin the private sector of the United States: make every private company in the United States operate under the Federal Acquisition Regulations and they'd be ruined because the flexibility is so limited and the bureaucracy is so much. Now is it ever going to change? I don't think so. If you have a democratic government and your stated objective is to treat everybody fairly, you just have taken on a whole load of rules and you've permitted vendors to claim that they weren't treated fairly and then it goes back and you have a whole year's review and this, that, and the other. So we move very slowly, very ponderously, very carefully because of wanting to operate in a democratic fashion in

government. So, it's just a general issue that, unless you're good at [it], unless you realize how the game has to be played, you can bumble things up badly. Now again, knock on wood, NSF has been pretty successful. Well, it should. We're relatively small, as federal agencies go, and have remained relatively agile, and I can't point to anything in our IT developments that I would say has been a major issue, like some of these other projects that we pointed to. Knock on wood; I hope I haven't jinxed the operation by saying that now.

Yost: Obviously, you were involved in all different kinds of oversight, as Chief Information Officer. How directly did you stay involved with what was going on with FastLane and with E-jacket?

Strawn: By the time I'd become CIO in '02 or '03, FastLane was pretty much in a maintenance standpoint, and I was more interested in keeping a close eye on things that were under development. So, I would say therefore, I had more involvement with E-jacket than I did with FastLane; very high level oversight with FastLane; medium to high level with E-jacket; was quite involved with some of the Research.gov stuff recently because I was a big booster of the data mining and other novel services that I thought we were in a position to do now, and so I inserted myself a little more into the process to try to get those things front and center in new requirements.

Yost: You've gone into this some already, but what ways could NSF use FastLane to more effectively advance the research enterprise?

Strawn: I think extensive data mining, both of proposals and ultimately, of the resultant research information, to do everything from facilitating finding reviewers to dividing proposals up modularly for panels; to understanding better the shape of the research activity of the community; to helping us tighten up our solicitations for proposals; a whole raft of things that have started and will sort of flower over this decade.

Yost: Looking back at FastLane, which is now over a decade old, are there lessons that NSF and other organizations could learn from FastLane to inform the development of cyber infrastructures?

Strawn: Absolutely. Software techniques for developing large program applications have progressed since the mid-90s. One of the phrases we use in the business is, "from tightly coupled modules to loosely coupled modules." Web services is a name that you hear that is floated around now. In the old methods, such as we used with FastLane, which were undoubtedly current when we used it, you end up with a big program, now all of a sudden if you change one of the modules in that program, you have no way of knowing for sure how many other modules get affected by it, in part, or maybe in most part, because of the way information is exchanged between the modules; by flat files, where you've got a whole bunch of information, you slap it out there, and this other module uses it but now you change it a little bit, well, what do you know, another module was using it too and now the information is in a different form. So that causes the maintenance costs associated with that type of a big software engine to be relatively high because when you

do make change — and you do have to make changes either with upgrades or just correcting bugs that you eventually discover — you introduce new problems that you then have to solve. It appears that the new, looser coupled technologies, which formally define the data interfaces between modules, rather than slap some data out there and let somebody else get it, should have a cost of lowering the maintenance, over the long haul, of big projects like that. We've at least begun using technologies in Research.gov and so I think we will find that maintenance of Research.gov will be a lower cost commodity because of more advanced software engineering techniques used to put this thing together.

Yost: What organizations, other than NSF, are involved in Research.gov?

Strawn: Our contractors, of course, and I think since we're using these modern techniques, we can make the services available; we call it a service-oriented architecture. You can make the services available on a piecemeal basis, so NASA has signed up for some services, DOD has signed up for some services, Department of Agriculture has signed up for some services. Those are three of the other agencies that I'm aware have signed up to use some of the services in Research.gov.

Yost: And finally, you've given us a number of possible avenues to investigate with our research. Are there other topics I haven't brought up; questions I haven't asked that would be useful to understanding the history of FastLane?

Strawn: No, you've been pretty thorough and I've been broadly ranging with my answers. I talked about the Pisani software. Pisani was Suzi Iacono's Division Director at the time this test software that made extensive use of data mining was done. He's, I think, Provost at a New Jersey institution right now. Suzi can probably tell you where he is. But he and his graduate students have written one or more papers on that subject and I would be pleased if you journeyed up to Pennsylvania or to New Jersey and found him and took his information about it. He's an information specialist and he did his best to add value to the FastLane milieu by those experiments. You would expect a researcher to be ahead of the contractors and his research, what he had us do, was something that was beyond what our contractors could implement, in my opinion. So when it came time to, could we commercialize that, could we put it into our stuff in a general way, is just a generation ahead of things. We'll be doing, in five years, the stuff that Pisani did for us five years ago. So that's a nice forward-looking piece of technology that I think would be nice if you included in your piece.

Yost: We'll definitely do that. Thank you very much for your time and your insights.

Strawn: Pleasure.