Testimony of Michael I. Shamos before the

Election Assistance Commission Technical Guidelines Development Committee Subcommittee on Computer Security and Transparency

September 20, 2004

My name is Michael Shamos. I have been a faculty member in the School of Computer Science at Carnegie Mellon University in Pittsburgh since 1975. I am also an attorney admitted to practice in Pennsylvania and before the United States Patent and Trademark Office. From 1980-2000 I was statutory examiner of electronic voting systems for both Pennsylvania and Texas and participated in every voting system examination held in those states during those 20 years. In all, I have examined over 100 different electronic voting systems, which I believe to be the largest number ever inspected by one person in the United States. I am teaching a course in Electronic Voting at Carnegie Mellon University this semester and am participating in a review now being conducted by the Commonwealth of Pennsylvania on the security of voting systems in all 67 Pennsylvania counties.

The Guidelines Development Process

Before we can even begin to discuss what evaluation guidelines ought to be, we have to discuss the process of developing guidelines in the first place. I want to start right away by arguing that the process I am involved in here in this room today is not the right one. Interested parties should not have to wait for an invitation to a hearing, then travel to the hearing at their own expense to explain their possibly complicated opinions in a few minutes via a formal statement. I've heard the claim that the process is open because there will be a defined period in which written public comment will be permitted. That resembles the federal regulatory process, which is well understood by lobbyists and special interest groups, but is hardly conducive to participation by the general public. And I hope that the IEEE model, in which draft standards are available on payment of a fee, is not adopted.

Such conditions do not promote a genuinely open process. It may be true that there a numerous guidelines in other fields in which the public has little or no interest, such as standards for cement to resist external sulfate attack[1]. Voting is a topic that affects every member of the public. It is a subject in which the public is not inclined to trust anyone, including politicians, government officials, computer scientists and standards makers. If they don't trust us, they must have a fluid way to participate in the process.

Once published, standards tend to become fossilized. It is only of minimal help to refer to such standards merely as "guidelines." As more manufacturers invest in designs and technologies that conform to given guidelines, they become highly resistant to allowing changes to be made, and certainly do not encourage them. This may account for

[1] Cliftin, J.R. et al., "Standards for Evaluating the Susceptibility of Cement-Based Materials to External Sulfate Attack." Available at http://fire.nist.gov/bfrlpubs/build99/art083.html. No disrespect is intended – the point is merely that a cement standard is unlikely to excite the interest of the general public.

the fact that the current Federal Voting System Standards do not address at all many computer security threats whose risks became known after the standards-making process was well underway. The result is that insecure systems routinely attain qualification when tested to these inadequate standards. To the extent that jurisdictions rely on federal qualification as a substitute for thorough certification, the FVSS have actually reduced the level of voting system security in the United States.

Guidelines must be responsive to development in the field. If new security exploits become known, the guidelines must be revised quickly to plug them. Otherwise, we will repeat the current situation, in which lawsuits are being brought to prevent the use of systems that were previously certified and might well continue to satisfy current certification criteria.

I wouldn't be complaining if I didn't have a proposed solution. I believe that the Internet model of Requests for Comment and the resulting development of guidelines through sustained participation by any knowledgeable and interested party, without barrier or formality, is the right one.

Will time-varying guidelines make it more difficult for vendors and jurisdictions to conform them? Yes, but that is the cost of keeping up with an ever-widespread and lever intruder community.

Scope of the Guidelines

Voting is a process that begins with registration and ends with death. At least, it should end with death but frequently does not because of the weakness of registration system in properly purging voters form the rolls. This is only one example in which external administrative factors influence the integrity of elections. Security guidelines must take into account the fact that voting is really encompasses over 10,000 processes supervised by tens of thousands of people and administered by over a million poll workers of varying levels of interest, trustworthiness and training. Having a completely secure voting machine is not sufficient if ineligible people are allowed to vote on it, or if eligible people can vote more than once.

Any examination of security should be comprehensive and should not focus exclusively on any one technology or method of voting. The claim is made, for example, that DRE machines are insecure, so therefore we should return to optical scan voting. But no one has done a security evaluation of optical scan systems, particularly with respect to physical custody of the paper ballots. These ballots spend long periods of time out of sight of the public and, in many cases, outside the sight of any election judge. They are sealed with plastic seals that are easily duplicated. If the Guidelines are to deal effectively with optical ballot, they must take ballot-handling into account.

Nature of the Guidelines

The FVSS are pass/fail in nature. Either a system meets a standard or it doesn't. Such a threshold system provides scant incentive to a vendor to exceed a standard, let alone expend additional funds to develop improvements. David Chaum, who will speak later today, suggests that the guidelines should not be binary, but should contain a rating mechanism so the quality of various systems can be evaluated along several dimensions. I am in wholehearted agreement with this idea because it allows comparison of different voting systems and methods, and provides an incentive for vendors to surpass one another.

There is much disagreement among experts as how good a system must be to meet a guideline. Some security specialists seem to suggest that a system must be perfect in order to be used in an election. That is, it can exhibit no vulnerabilities whatsoever. If that is indeed the test, then we can all go home, for the task will be futile. Others propose that since voting is a matter of national security, the same level of vigilance must be maintained as we use to maintain operation battle plans or nuclear launch codes. That may be so, in which case the assumption should be stated explicitly and we should be prepared to bear the expense of containing top-secret clearances for 1.4 million poll workers many of whom, it turns out, have criminal records, which is not, in most states, a barrier to handling election materials.

Audit Trails and Their Role in Security

There has been much written and spoken about audit trails that is incorrect and that should be corrected because of the important role audit trails play in security. First we must draw a distinction between an audit, which is the process of verifying that no irregular events have taken place during the election, such as resetting the machine counters during voting, or counting a precinct twice, and ballot reconstruction, which is retrieving the individual ballot images of each voter who has voted, usually for recounting. Let us call this latter a "ballot trail" and the former an "audit trail." We need separate terms for them since they are so often confused.

The argument is made that a ballot trail must be on paper since no electronic trail can be trusted. The argument is wrong in several respects. First, whether or not a ballot trail can be relied upon depends on whether it was created correctly and whether it has been preserved without tampering. While voter verification of a ballot trail may indicate that it has been created properly, it affords no assurance that its integrity will be maintained after the voter has left the polling place.

Second, an electronic ballot trail is fully reliable if it is well-designed, can be tested and is impervious to attack. On the other hand, if the ballot trail, paper or electronic, has been compromised, it is of no value except that it may provide evidence to forensic investigators of how the compromise was accomplished. The widespread movement toward paper ballot trails has essentially foreclosed necessary research into alternatives.

Some computer scientists have alleged, without providing even a convincing argument, let alone a demonstration, that DRE machines do not allow a "meaningful recount." Certainly if the trail mechanism has been infiltrated, the trail will not permit a recount, but this is also true of paper. A functioning electronic ballot trail mechanism is more reliable than any paper one can possibly be.

Despite vendor claims to the contrary and the evident belief of some election officials, there presently exists no commercial implementation of a voter-verifiable paper trail. There are paper-trail machines, to be sure, but each of them inserts bar codes, numbers or cryptographic indicia onto the ballot that the voter cannot decipher or understand, let alone verify. These indicia, supposedly used to prevent introduction of spurious ballots, can in fact be used to invalidate perfectly proper one, and the voter will never be any the wiser.

Some current commercial paper trail implementations have such severe flaws that they violate the laws of the very states in which they have been so hastily adopted. The Sequoia system used in Nevada two weeks ago maintained a consecutive reel-to-reel

paper trail. This means that the ballot of the first voter was first on the tape, the ballot of the last voter was last on the tape, and all the rest were recorded in sequential order in between. This is a complete violation of ballot privacy, since anyone with access to the tape and a poll list could reconstruct the vote of every voter.

The Nevada Revised Statutes state, at Section 293B.065, that the "voting system must secure to the voter privacy and independence in the act of voting." Possibly this was interpreted by Nevada officials mean that you have privacy during the actual act of voting, but after you leave the polling place it's okay for poll workers to review how everyone voted.

Audit trails, as distinguished from ballot rails, are a different matter. They normally maintain, when properly implemented and not compromised, a record of events surrounding the election, such as the opening of polls, loading of ballot styles, recording a vote, performing administrative functions, etc. These are necessary to reconstruct the steps that were taken during the election so that procedural integrity can be assured.

Recounts

The belief is widespread that a ballot trail can be used to accomplish a reliable recount. Aside from recounts that are mandated by law, such as the 1% manual recount in California, recounts generally occur only when balloting has been close. Where the recount involves physical ballots, the people who handle them during the counting have the ability to alter them, dispose of them or substitute others. This is a time-honored method. Because voting has been close, only a small amount of manipulation is needed to affect the outcome. By contrast, properly encrypted electronic records cannot be lost or altered.

Security Evaluation

Security cannot be evaluated without a well-articulated threat model. The reason is that without such a model it is impossible to determine whether the system adequately resists the threats. The model is just as important as the standards that relate to it. Unless the model is comprehensive and agreed upon, any resulting standards will induce a false sense of security even if they are met.

Let's be serious. No set of static equipment standards will do the job. Every system is used within a real-life context of state, county and polling place culture, laws and procedures. Election officials, inspectors, watchers, voters and politicians all operate with differing degrees of vigilance and experience. The notion that some collection of standards will suffice to ensure voting system security is a pipe dream.

Voting system must be evaluated *in situ*. Whether a system meets the guidelines cannot be determined by sending machines to a laboratory for a period of time. The procedures, policies, safeguards that are actually used must be evaluated. It is conceivable that a given system will be secure when used in County A but not in County B. A laboratory test will never reveal this, and any mechanism that purports to establish a certified list of testing laboratories will not solve the problem. It is of little consequence to pronounce that a given version of firmware is correct or secure if the voting machines without controls on access to the machines in which the firmware is installed.

This is not to say that laboratories have no role, but that role must be carefully understood and circumscribed. Currently the FVSS afford few procedures for the

laboratories to follow in testing voting systems. This means that even if the labs are completely diligent, when they certify that a system has met the standards the public has no way to find out what the lab did to verify that fact, and what deficiencies were observed. Aside from being opaque, the process is also inadequate because it does no include observations of use of the system in practice and does not include "red team" exercises designed to uncover serious flaws that may not technically violate the standards.

To summarize, my most important recommendations are that the security of a voting system must be evaluated holistically as an entire system rather than as set of hardware and software components. The process for developing guidelines must be as open and inclusive as possible and that the guidelines should deal with quantitative measures of performance, not simply with pass/fail requirements. I thank you for the opportunity to address the Subcommittee here today.

Biography of Michael I. Shamos

Michael I. Shamos is Distinguished Career Professor in the School of Computer Science at Carnegie Mellon University, where he serves as Co-Director of the Institute for eCommerce and a Director of the Center for Privacy Technology. He has been associated with Carnegie Mellon since 1975. He is Editor-in-Chief of the *Journal of Privacy Technology*.

Dr. Shamos received an A.B. in Physics from Princeton University, an M.A. in Physics from Vassar College, M.S. degrees from American University in Technology of Management and Yale University in Computers Science, the M.Phil. and Ph.D. in Computer Science from Yale University and a J.D. from Duquesne University. He is a member of the bar of Pennsylvania and the United States Patent and Trademark Office.

From 1980-2000 he was statutory examiner of computerized voting systems for the Secretary of the Commonwealth of Pennsylvania. From 1987-2000 he was the Designee of the Attorney General of Texas for electronic voting certification. During that time he participated in every electronic voting examination conducted in those two states, involving over 100 different voting systems accounting for more than 11% of the popular vote of the United States in the 2000 election.

Dr. Shamos has been an expert witness in two recent lawsuits involving electronic voting: *Wexler v. Lepore* in Florida and *Benavidez v. Shelley* in California. He was the author in 1993 of "Electronic Voting — Evaluating the Threat" and in 2004 of "Paper v. Electronic Voting Records — An Assessment," both of which were presented at the ACM Conference on Computers, Freedom & Privacy. He has provided testimony on electronic voting to the Pennsylvania legislature and the U.S. House of Representatives Science Committee and the House Administration Committee. He was a member of the SERVE Project Review Group and the recent National Research Council Workshop on Electronic Voting.

Dr. Shamos has not received any federal grants or contracts during the past fiscal year.

Further information is available at http://euro.ecom.cmu.edu/shamos.html.