Testimony of Dr. Dan S. Wallach
NIST/EAC Technical Guidelines Development Committee
September 20, 2004

Thank you very much for holding today's hearings. I appreciate the opportunity to speak to you today about the security issues in the electronic voting systems that are increasingly being used in the U.S.

I am an assistant professor in the department of computer science at Rice University in Houston, Texas. I earned my bachelor's degree at the University of California, Berkeley in 1993 and my doctorate degree at Princeton University in 1999. I study computer security and have published over forty refereed academic papers on computer security and related topics[1]. I have investigated the security of web browsers and web servers, looking at how to keep your computer from being hijacked just because you clicked the wrong link. Several of my contributions shipped as part of Netscape's Communicator in 1996 and are now part of every Java system in use today. I have also investigated the security of other networked and distributed systems. In general, I look at computer security as an engineering problem. The goal in designing and building a secure system is to understand the threats the system might face and to build in appropriate safeguards to protect against those threats.

I first began examining electronic voting systems in 2001 when I was invited to testify before the Houston City Council about the Hart InterCivic eSlate voting systems that were being adopted by Harris County. Last summer, I co-authored a report with Adam Stubblefield, Tadayoshi Kohno, and Aviel Rubin, at Johns Hopkins University, that examined the design of the Diebold AccuVote-TS voting system; that paper appeared recently at an IEEE security conference[2]. I have also co-authored the "frequently asked questions" document for VerifiedVoting.org[3] and conducted research on the ability for testing authorities to detect flaws in voting systems. Based on my research, I have come to conclusion that paperless electronic voting systems (also called "direct recording electronic" or "DRE" systems) are fundamentally insecure and do not provide sufficient protections against the sorts of fraudulent behavior that have been historically taken to manipulate the outcomes of elections in the U.S.

Threat Models

When considering the security of any computer system, whether for voting or for other applications, the analysis always starts by looking at the threats the system will face. Threats can include everything from loss of electrical power or other physical issues including dropping the machines on the floor. Threats might include software bugs or mistakes in the machine's configuration and installation. When these things have happened in the past, the results have often been inexplicable, casting serious doubts on the validity of many elections. For example:

> Florida's official line is that its machines are so carefully tested, nothing can go wrong. But things already have gone wrong. In a January election in Palm Beach and Broward Counties, the victory margin was 12 votes, but the machines recorded more than 130 blank ballots. It is simply not believable that 130 people showed up to cast a nonvote, in an election with only one race on the ballot. The runner-up wanted a recount, but since the machines do not produce a paper record, there was nothing to recount.
>
> In 2002, in the primary race for governor between Janet Reno and Bill McBride, electronic voting problems were so widespread they cast doubt on the outcome. Many Miami-Dade County votes were not counted on election night because machines were shut down improperly. One precinct with over 1,000 eligible voters recorded no votes, despite a 33 percent turnout statewide. Election workers spent days hunting for lost votes, while Floridians waited, in an uncomfortable replay of 2000, to see whether Mr. McBride's victory margin, which had dwindled to less than 10,000, would hold up.

---

[1] `http://www.cs.rice.edu/~dwallach/pubs.html`

[2] Tadayoshi Kohno, Adam Stubblefield, Aviel D. Rubin, Dan S. Wallach, *Analysis of an Electronic Voting System*, 2004 IEEE Symposium on Security and Privacy (Oakland, California), May 2004. Also available online. `http://avirubin.com/vote/`

[3] `http://www.VerifiedVoting.org/drefaq.asp`

> – "Florida as the Next Florida," *New York Times* (unsigned editorial), March 14, 2004.[4]

An important class of threat that has not been considered as carefully as it should be is the threat of *software tampering*. At any point in a voting machine's life, from the manufacturer's shipping dock through intermediate storage to the day of the election, a voting machine could potentially be *reprogrammed* to report incorrect results. Such "Trojan Horse" attacks have already occurred in the gambling industry. For example, Ron Harris, a former member of Nevada's Gaming Control Board, was convicted in 1998 for tampering with computerized slot machines.

> Harris inserted a computer program into a device used by control board employees to check the proper functioning of slot machines. When the testing device was used by control board employees, it downloaded a cheating program, called a gaff, into computer chips in the machines.
>
> Harris then recruited a trio of friends … to play slot machines rigged with the cheating program. Inserting a specific series of coin bets allowed the program to take effect and award jackpots.
>
> – "Former gaming official sent to jail for slot scam," *Las Vegas Review Journal*, January 10, 1998.[5]

Fundamentally, this is very little difference between an electronic gambling machine and a paperless voting machine. Somebody with private access to our electronic voting systems could perhaps arrange for the installation of a modified software in the same manner as Harris corrupted some of Nevada's gambling machines; when a co-conspirator performs an unusual write-in vote or otherwise makes an unlikely series of button presses, the voting machine might change its records to artificially favor one candidate over another. Despite the precautions we might take, any paperless DRE voting system will be vulnerable to this class of threats.

Mitigating Strategies

The states of Ohio and Maryland, partly in response to our report on the insecurity of the Diebold voting system, commissioned independent studies to either confirm or refute our findings, where we showed how normal voters could cast multiple votes using "homebrew" smartcards; we showed how Diebold's incorrect use of cryptography would allow the voting records to be silently modified; and we showed how Diebold's software engineering discipline was far below the standards that would be applicable in other contexts. These reports, from SAIC, RABA, Compuware, and InfoSentry, generally confirmed our technical findings or, in the case of the Compuware report, were not able to reproduce them but did not rule them out. These reports and our report generally disagree on the impact of these technical findings and what strategies may be necessary to adequately mitigate these serious technical flaws. Our position is that, regardless of whether the software in the Diebold or other voting machines is improved to better resist attacks, bugs will *always* occur and the risk of tampering cannot be overcome. In particular, we believe that while "logic-and-accuracy testing" can sometimes detect flaws, it will never be comprehensive; important flaws will always escape any amount of testing. Likewise, the certification process and the efforts of independent testing authorities (ITAs) such as Wyle Laboratories are insufficient to demonstrate, beyond a doubt, that these voting machines will operate properly.

An important and unanswered question is whether any ITA can ever apply sufficient scrutiny to the voting machines' software to truly detect whether that software operates correctly in all circumstances. As an exercise in my graduate-level computer security class last fall, we asked the students to first take on the role of a corrupt software developer trying to hide subtle but significant flaws in the software of a voting system that we had already built in-house. We then swapped their work with other students, who were asked to audit the code, looking for flaws created by the first group. Our study showed that, while a number of flaws were discovered,

[4] http://www.nytimes.com/2004/03/14/opinion/14SUN1.html

[5] http://www.reviewjournal.com/lvrj_home/1998/Jan-10-Sat-1998/news/6745681.html

many subtle and clever flaws passed our internal audits. We believe that ITAs will be unlikely to do much better. Our results have been published in an IEEE magazine.[6]

Other mitigating strategies have been proposed, including tamper-resistant measures such serially numbered locks or tape that changes color if somebody attempts to remove it in order to access the voting machine's internals. While these measures are well intentioned, we believe that a sufficiently motivated adversary can either tamper with the machine before these tamper-resistant measures are taken, or can fabricate his or her own locks or tape to reinstall on the machine after any tampering has taken place. Consider both Harris and Travis Counties, among others, where we are relying on exactly this technique to protect our voting machines. I recently gave a talk where I mentioned a "bad-guy" who might break into the high school cafeteria at night to tamper with the machines. A member of the audience, a local poll worker, said that I had it all wrong; he took the voting machines home with him at night. Really! The bar for becoming a poll worker is generally not very high. Travis County, for example, only requires that they not be convicted felons. How reassuring. We're allowing these people to take voting machines into their homes, with their garage full of tools, and we're claiming that we can somehow physically keep them from prying machines open, modifying their internals, and then sealing them back up again. Such claims are very difficult to support. We need to do better.

Another proposed mitigating strategy is the use of *open source software*, that is, making the software source code for the voting machines available for anyone in the public to read and examine. For example, Australia developed a voting system called eVACS (electronic voting and counting system). The Australian government contracted with a private firm to develop the software, which is currently available to be freely used by anybody else, anywhere in the world, at no cost.[7] Such an open source election system would allow any interested third-party to make its own examination of the security of a voting system, possibly finding security flaws and bringing them to the attention of the system's developers. Most voting system vendors, however, consider their source code to be a trade secret. We were only able to analyze Diebold's system as a result of their inadvertent release of their source code on the Internet. This demonstrates that any security protection that might be gained from keeping the code private is temporary, at best. It's better for a system to be designed to be secure *regardless* of what knowledge is possessed by a would-be attacker. And, because open source software gives independent third parties the ability to make independent evaluations of the integrity of an election, open source software increases the *transparency* of the election, which can clearly help increase voter confidence in an election's outcome. Unfortunately, even if the source code is public, subtle but exploitable flaws may still persist in it for years. Likewise, we have no easy way to guarantee that a voting machine is running the "official" software version rather than a maliciously modified version. Open source code is valuable for an election's transparency, but it is not sufficient to make any security guarantees.

(An important confusion about open source code is what sort of license might be used for the open software. For security purposes, it is only interesting that the software be readable by anybody. There is no requirement that the software be given away to be used by third-parties at no cost. As such, a vendor may still keep a proprietary interest in its code while still publishing its code. Standard copyright and patent laws would protect the intellectual property of the vendor.)

The most robust mitigating strategy of which we are aware is the use of a *voter-verifiable audit trail* (VVAT). Most commonly, a VVAT system is a normal DRE voting system with an attached ballot printer. Voters can see and verify their ballots, but cannot keep them. The ballots are stored in traditional ballot boxes and tabulated at the end of the election. The security benefits of such a system are easy to understand. If the voting machine malfunctions, either as a result of a software bug or as a result of deliberately software tampering, then the printed paper ballot would be incorrect; the voter, after inspecting the paper ballot, would reject it. This would create a "spoiled ballot," for which well-understood procedures already exist to destroy the spoiled ballot and give the voter another opportunity to cast his or her ballot. In a VVAT system, *the correctness of the software no longer matters*. Either it consistently produces paper ballots that match voters' intent, or it is taken out of service.

---

[6] Jonathan Bannet, David W. Price, Algis Rudys, Justin Singer, Dan S. Wallach, *Hack-a-Vote: Demonstrating Security Issues with Electronic Voting Systems*, IEEE Security & Privacy Magazine, volume 2, number 1, January/February 2004, pp. 32-37. Also reprinted by *ComputerUser*, March 2004. Also available online. `http://www.cs.rice.edu/~dwallach/pub/hackavote2004.pdf`

[7] `http://www.softimp.com.au/evacs.html`

An important benefit of VVAT over paperless DRE systems is the ability to audit the election. VVAT paper ballots are collected and stored in traditional ballot boxes such that they can be counted to determine the final election tallies. Because they were printed by computers, they can be read by other computers using optical character recognition (OCR) tools. They can likewise be read by humans, if for whatever reason the electronic counts are considered unreliable. The VVAT ballots may also contain cryptographic security measures, perhaps printed as a bar-code, to provide protection against ballot stuffing attacks.

VVAT ballots can perhaps best be considered to be a strong form of *evidence* of the voter's intent. While numerous different procedures can be imagined for how and when that evidence is considered, whether as part of vote tallying or as part of post-election audits, nonetheless the evidence remains, and traditional legal notions of the chain of custody of that evidence can be applied to maintaining it. Because this evidence was seen (and verified) directly by the voter and could no longer be subject to electronic tampering, it provides an important hedge against any of a variety of failures that might occur in the electronic domain.

Certainly, the notion of having independent printed records of important data is not an idea unique to voting. Our banking industry, despite all of their computers, generates huge amounts of paper. Every ATM prints a receipt for its transactions. Credit card transactions likewise generate paper receipts. Furthermore, banks send every customer a printed statement at the end of the month. The existence of these redundant records allows for inconsistencies and fraud, which occur on a regular basis, to be detected and corrected. VVAT provides this same level of assurance to our election systems.

Criticisms of Voter-Verifiable Audit Trail Systems

The concept of VVAT systems have been scrutinized in a number of venues, resulting in many common criticisms that I would like to discuss.

Claim: VVAT printers will jam and require costly maintenance.

Day in and day out, cash registers, ATMs, and numerous other machines print receipts without requiring any maintenance. If VVAT technology is adopted, industrial-grade printers can be specified that will be more than sufficient for election duties. Pre-election testing and maintenance can determine whether the printers are working properly. And, in the worst case, printers can be designed to be easily removed and replaced, in the field, during an election.

Claim: VVAT systems will cost more money to add the printers and maintain the paper ballots.

While printers may add some cost to DRE voting systems, they will ultimately save money in a number of ways. When a county or state buys an electronic voting solution today, they buy everything from a single vendor to guarantee the machines interoperate correctly. In a VVAT system, a county or state could mix and match vendors, so long as the exact format of the paper ballot (i.e., fonts, line spacing, margins, and so forth) is standardized. This would allow different vendors to sell the ballot preparation and the ballot tabulating systems, increasing competitive pressures and reducing costs. Furthermore, VVAT systems do not require the chain of custody of the voting systems to be carefully maintained to prevent tampering. Either a VVAT system presents the correct printed ballot to the voter, or it is pulled out of service.

While many election officials would like to eliminate the burden of warehousing and otherwise managing ballot boxes with paper ballots, this is a necessary cost to protect the auditability and integrity of the election.

Claim: VVAT systems do not satisfy ADA or HAVA requirements for accessibility.

VVAT systems have the same accessibility properties as paperless DRE systems. They can support headphone jacks and large text for blind and low-vision voters. They can support multiple languages. They can present a "review" screen with all of the voter's selections displayed. They can eliminate overvoting, can warn voters if they undervote, and can support other desirable features such as straight-party voting, instant runoff voting, or other non-traditional election styles. While a blind voter may not be able to read the VVAT paper ballot, *the voting machine cannot distinguish a blind voter from a sighted voter*. Just as blind people use ATMs and can trust

they will receive the correct amount of cash, they can similarly trust that VVAT systems will not be able to discriminate against them.

Claim: VVAT systems rely on paper, which has its own long history of fraud.

A VVAT printer is significantly different from punch cards or optical scan machines. Traditional ballot-stuffing attacks can be defeated by having the VVAT systems apply cryptographic digital signatures to the paper ballots, perhaps printed as a bar-code. Likewise, ballot "serial numbers" could be encoded on the paper ballot and in electronic records maintained within the computer. These records could later be reconciled to make sure the electronic and paper records agree with one another. In the event that paper records exist without electronic equivalents, then procedures would be necessary to determine how the electronic records were lost and to verify the serial numbers and digital signatures on the paper ballots. Likewise, if paper ballots are lost, then electronic records from the voting machines could be used as a backup. In general, when discrepancies occur, the paper ballots should be considered to be the primary record of a voter's intent because the voter actually *saw* the paper, while the voter did not see the bits inside the computer.

Claim: VVAT systems will be slow to generate election results.

VVAT systems, because they are built using computers, can certainly keep electronic tallies, and these electronic tallies can be rapidly tabulated. Such early tabulations should be considered to be as accurate as early returns or exit polls. They do reflect the will of the electorate, but they should not be certified until the paper ballots have been scanned, tabulated, and reconciled against the electronic records.

Claim: VVAT systems will be difficult for relatively untrained poll-workers to manage.

A VVAT system is comparable to current DRE systems, in terms of manageability. Traditional paper-based systems, particularly optical scan voting systems, are significantly simpler to set up and to explain to both poll workers and normal voters.

Conclusions

In our analysis of DRE voting systems, including the Diebold AccuVote-TS, we have found significant security vulnerabilities that could call into question the integrity of an election's results. In the event of significant tampering with the machines' software, insufficient evidence will remain to determine which, if any, machines had been tampered with and what damage may have been done to the election results. While computer technologies can provide significant human-factors and accessibility benefits, these benefits are meaningless if the election is vulnerable to significant fraudulent activity. As a result, we believe that paper ballots that can be read and verified by voters (a voter-verifiable audit trail), must be an integral part of modern elections.

My recommendations are as follows:

- The TGDC should develop and promulgate "best practices" for mitigating the security risks inherent in the present generation of DRE systems. In particular, we should require that a voting machine, a vote tabulation system, or any sort of vote storage memory card, is never held in the custody of any single person, at any time.

- The TGDC should recommend the decertification of the present generation of DRE voting systems until more secure replacements have been developed. To make a seamless transition, a grace period should be introduced, perhaps similar to the cut-over schedule now mandated in California.

- The TGDC should specify far more stringent standards concerning the security and tamper-resistance of voting systems. Rather than mandating a specific technology, such as a voter-verifiable audit trail, the TGDC should specify stronger notions of auditability and voter-verifiability than exist in present statues or FEC standards. This will allow for future technologies that may not yet be commercially viable, such as novel cryptographic approaches, to be applied in the future. Likewise, experts have proposed a number of specific changes that could be applied to the current generation of DRE systems such as having these machines display a "cryptographic hash code" of their software when they first boots; done properly, this would increase the difficulty of tampering with a machine's software. The TGDC needs to look at both short-term fixes and long-term strategies to address flaws in current voting systems.

- The TGDC should carefully reexamine the national and state procedures used to certify voting systems. Consider the state of Texas as a reasonable common example. The secretary of state has an advisory board that issues recommendations on which voting systems do and do not satisfy Texas requirements. Presently, this board only performs a superficial analysis of DRE voting systems. None of them read the source code to these systems, despite being provided with this code; they have neither the time nor the expertise. Instead, this board relies on "independent testing authorities" (ITAs) to perform such an analysis, but that analysis appears to also be quite superficial. Instead, the TGDC should recommend that states, or perhaps the ITAs themselves, contract with outside expert organizations to perform so-called "tiger team" or "red team" analyses of any voting system to be considered. Such analyses, performed on behalf of the states of Maryland and Ohio, have found significant issues in every voting system they have ever examined. While companies should be required to pay for these analyses, the states should be able to choose the 3rd party expert examiners to guarantee that companies do not simply shop around for analysts willing to perform the minimal and simplistic tests that are currently being performed. I also strongly support the proposal, echoed in other testimony, that the results of such analyses should not be a simple "thumbs up" but should rather be a series of numbers, along several different axes, that describe how well the system performs. This will give incentives to vendors to go beyond current standards.

- Today, every voting system requires separate certification, both at the federal and state levels. This certification can be quite expensive and slow, and more stringent standards and testing will only exacerbate this problem. To address this, open source software can and should have a significant role in the security of our nation's election systems, including front-end voting terminals and back-end vote tabulation systems. By separating the development of the software from the specification of the voting hardware, states can collect bids from "clone" hardware makers who will not need to perform these expensive software certifications. When a state mandates a software standard, the state would perform the certification of its own software. Furthermore, multiple states can share the burden of developing these software systems, sharing a common core with customizations to meet election statues and customs that vary from state to state. Given the variety of statues and customs from state to state, it would be unlikely that a single system could be developed nationally to satisfy the needs of all states. By encouraging individual states to pursue the development of such systems, perhaps with grants to support such developments, states will ultimately save money and improve the quality of their election systems.