Testimony of Prof. Dan Wallach
Tennessee Advisory Commission on Intergovernmental Relations (TACIR)
September 27, 2007

Chairman Rinks, Mayor Rowland, members of the committee, it's my honor to speak to you today about security, privacy, and related issues with the electronic voting systems used here in Tennessee and across our nation.

I am an associate professor in the Department of Computer Science at Rice University in Houston, Texas and the associate director of NSF's ACCURATE (A Center for Correct, Usable, Reliable, Auditable and Transparent Elections). A collaborative project involving six institutions, ACCURATE is investigating how security technologies may best be applied to electronic voting systems. I have also served as an expert witness in seven different lawsuits concerning electronic voting, including the recent Congressional election controversy in Sarasota, Florida.

I first began examining electronic voting systems when I was invited to testify about them before the Houston City Council in 2001 concerning the Hart InterCivic eSlate systems then being adopted in Harris County. Since then, I have co-authored several research papers on the topic, including a landmark study in 2003 where we performed a source code security analysis of the Diebold AccuVote-TS voting system. This summer, under the auspices of the California Secretary of State's "Top to Bottom Review", I was part of a team that performed a similar analysis of the source code of Hart InterCivic's systems. Other teams considered the security of voting systems from Diebold and Sequoia, as well as accessibility and documentation issues. The California study was the most comprehensive study of its kind ever performed, and I'm pleased to report that the State of Ohio is about to begin a follow-on study.

What we found was troubling. These voting systems didn't seem to have been designed or implemented with security in mind. At best, security appeared to have been bolted on as an afterthought. These systems failed to follow standard security design principles. The code was riddled with security holes.

We found serious security vulnerabilities in all three systems. We discovered that a lone individual, with no special access, could introduce malicious code into each vendor's voting machine. Once one machine is compromised, we discovered that the compromise could spread virally from machine to machine. When voting machines are returned to county headquarters on election evening to count the votes, the virus could travel alongside the votes and infect the county's election management computers. The virus could then spread to every voting machine in the county in the next election.

The viral vulnerabilities are troubling because they could be carried out by a single voter and could affect an entire county. Likewise, an attacker could introduce a virus by tampering with a voting machine while it is stored unattended overnight. For some of the systems, team members even demonstrated that a voter could introduce malicious code into a voting machine in under a minute, while in the process of voting, leaving poll

workers none the wiser. These flaws were systemic and surprisingly similar across all three voting systems we examined.

Many of these flaws were remarkable in how basic and profound they were. One vendor uses an identical cryptographic key, hard-coded into the software, in every machine they make. If every car in the parking lot had the same key, nobody would consider that to be a secure design! The same issue applies here. A different vendor used their own name as a hard-coded password – a basic, obvious design error. Another vendor's system allows an attacker to connect to a port on the back of the machine and issue commands that can read and write any address in the machine's memory. With this an attacker can extract secrets from the machine, replace its software with something malicious, or even overwrite previous votes. One machine's audio port, used for accessibility purposes and chatting away whether or not the voter is wearing the headphones, inadvertently broadcasts that audio as an AM radio signal that can be easily received at a safe distance. Our public reports were carefully written to convey the depth of the problem without providing a road map to prospective attackers.

Appropriately, California's Secretary of State took a variety of actions to limit the use of these systems, most notably allowing only one electronic voting machine per precinct in counties using Diebold and Sequoia systems. I would recommend similar precautions for all other electronic voting systems, including those not considered as part of the California study.

Perhaps the most important lesson we have learned from this analysis is that the existing regulatory framework has been insufficient to yield good voting systems. The "independent testing authorities" and the various state-level examination and certification procedures have demonstrably failed to previously identify flaws that we discovered with only a month of work, including the effort of crafting functioning exploits to demonstrate several of these flaws. This raises several serious concerns for how we must proceed. We could impose more stringent standards, but these machines arguably don't even meet the existing standards. We could likewise mandate better testing, effectively forcing vendors to reengineer their products from scratch. Maybe they will do better, maybe they won't. We could drop these machines altogether, moving to simpler paper ballot technologies, but those have many of their own security issues.

Any path, going forward, needs to recognize and distinguish between short-term procedural changes and long-term redesign. We cannot simply impose a greater burden on our poll workers to catch malicious voters "in the act." Instead, in the short term, we should focus on technologies that limit the damage that can be caused by a small number of malicious actors in the system, and today that means paper ballots, whether marked by hand or by accessible "ballot marking devices." Those ballots can be mechanically tabulated, but there must be a carefully crafted statistical auditing process to cross-check the electronic system against the paper for errors. Issues such as these are considered in Congress's H.R. 811 (the "Holt Bill"), which I strongly support.

In the longer term, we cannot continue with the current model where vendors have little effective oversight and produce products of insufficient quality. A variety of alternatives are available that can improve on the status quo. We could involve external experts earlier in the vendors' design processes. We could require vendors to publish their design documents and source code, as trade secrecy plays no useful role in protecting voting systems against attack. Voting systems can and must be engineered to resist attack from attackers who understand at least as much about their inner workings as we were able to learn in the California study.

At the end of the day, the purpose of an election is not only to name a winner, but more importantly to convince the loser that he or she genuinely lost the election. By virtue of their inadequate engineering, today's electronic voting systems provide ample room for doubt. When election results are tight, these issues become even more prominent. In the Sarasota case, now nearly a year after the election, the legal battle is still ongoing. Will the same issue occur here in a future Tennessee election? Perhaps, but with appropriate regulations, similar disasters can hopefully be avoided.