CERIAS Tech Report 2007-65

Voter Assurance

by Eugene H. Spafford

Center for Education and Research in
Information Assurance and Security,
Purdue University, West Lafayette, IN 47907-2086

National Academy of Engineering Website
http://www.nae.edu/nae/bridgecom.nsf/weblinks/MKEZ-744M69?OpenDocument

Voter Assurance

Eugene H. Spafford

For the good of the country, we need voting technologies produced and operated according to the best methods.

Voting with assurance seems to be an obvious, simple concept. A registered voter should be able to cast his or her ballot with the confidence that the vote will be counted as cast. Traditionally, paper ballots have seemed like a simple, comfortable voting solution. However, paper ballots in some forms can be easily manipulated, result in ambiguous interpretations (e.g., “hanging chads”), are sometimes error-prone, and do not provide a quick tally. In our technology-saturated society, we want results right away, and it would seem that technology could speed up vote counting and make it more accurate. Computers are being integrated into every aspect of our lives, so why can’t they work for voting, too? If we can use computers to control airplanes, factories, and ATM machines, we should certainly be able to use them in voting!

For a number of years, computerized voting machines have been used in various jurisdictions around the United States and around the world. As a result of the Help America Vote Act (HAVA), which was passed in 2002, thousands of new machines were purchased and deployed around the United States to replace paper-based and lever machines. The vendors of these machines repeatedly stated that they were accurate and secure, and election officials echoed their claims. Yet there have been repeated incidents of surprising vote totals, interface glitches, and unexplained behavior in some machines. As a result, voter confidence seems to be waning.

“Assurance” can be defined as a positive declaration intended to give confidence. In this context, it is a promise to voters that their votes will be counted accurately and fairly, and thus will figure in the overall outcome. Assurance means providing certainty about something, in this case, the certainty that each vote matters. This is essential to our republic, which depends on citizens voting to express collective choices.

We want people to vote. We want them to participate in our democracy and to voice their opinions. Thus we want them to have confidence in the honesty of the process, and we want them to have justified confidence that their votes will actually be counted.

A Lack of Confidence
A number of indicators show that some of the voting population does not have confidence, or is losing confidence, in the technologies that have been deployed as a result of HAVA. A study by the Pew Internet Trust found that groups of minority voters in the South do not trust the technology and believe it has been manipulated, or could be manipulated, so that their votes would not count. They expressed a significantly higher rate of distrust in computerized voting technology than the general voting population (Kohut, 2006).

A study by the University of Maryland Center for American Politics and Citizenship found that one of every 10 voters leaving the polling place did not have confidence in the technology they had just used—direct-recording electronic touch-screen voting machines (DREs) (UMd, 2002). Ten percent of voters, a significant number, were not certain their votes would be counted accurately. Add to this another large group that may not have bothered to vote because they also distrusted DREs.

Some have speculated that the significant increase in absentee balloting in recent years is at least partly attributable to voter distrust in electronic voting machines. For example, in California, 43 percent of the electorate cast their votes in the most recent election by absentee paper ballots, which may have reflected a lack of faith in electronic voting machines (Elias, 2007). These machines, which have been deployed throughout California, did have some problems, such as miscounts and lost votes, and public debate has arisen about whether certain machines were reliable and about problems with the systems.

Other states have also noticed an increase in the number of absentee ballots. Some computer scientists

and voting advocates have stated publicly that voters should use absentee ballots to ensure that their votes are counted if their local precincts use DRE machines without paper audit trails (note, that mail-in ballots may be even more susceptible to coercion and fraud). Coupled with voter unease, this may explain some of the increase in absentee balloting.

Many professionals working in the area of information security and reliability question some of the assumptions about voting technology. They have also expressed considerable doubts about pronouncements by a few people who claim to be experts in this area, particularly about how IT systems can be compromised, where failures can occur, and how such failures occur. Rather than reassuring real experts, these statements raise concerns because they are often simplistic or do not take into account known threat models.

Even some voting machine vendors admit that there is growing mistrust and, perhaps, for good reason. Dennis Medura, CEO of Accupoll, a now-bankrupt voting machine manufacturer, made the following statement about voting system software: “I do not feel that any of the vendors has a system that voters can trust” (Greene, 2006). He went on to say that vendors have misrepresented the robustness, stability, and security of their systems. This has long been suspected by many who have investigated these systems and, in some cases, had hands-on experience with the equipment. These experts can testify that the software and the hardware are not as well protected as vendors suggest.

Our republic depends on citizens voting to express collective choices.

Seeking Assurance
How do we go about instilling confidence that a system works the way it is supposed to? Absolute assurance, by its very nature, can never be achieved because we cannot meet every requirement of every individual. At the very least, some otherwise rational individuals are deeply suspicious, and meeting the conditions to satisfy them would be cost-prohibitive.

In addition, voters have many different requirements. Consider the requirements for ensuring access for voters with visual and hearing disabilities, limited motor skills and cognitive impairments, and for people unable to travel to the polls. Each of these groups needs different system interfaces that have dramatically different security and reliability issues.

Even more challenging, complete assurance means more than providing security and eliminating flaws. It requires that voters be able to understand the security properties and failure modes of a system well enough to have “observed confidence” that their votes will be counted. But because voters have varied backgrounds, some may not be able to judge the overall security or reliability of a complex software system.

The stakes in an election matter at almost every level.

Sources of Trust
One way to convince ourselves that something works is to examine and test it. Before we buy something, we examine it (or an exemplar) or try it in actual operation, look at it carefully, and subject it to tests to assure ourselves it will perform as expected. However, most people do not have the expertise, time, or resources to perform exhaustive tests on all aspects of every purchase.

An alternative is to find a trusted source that can verify and vouch for the results of testing. For example, we trust the judgments of Underwriters Laboratories because they publish their standards and we are familiar with their long history of certifying items. We may also trust governmental entities because of their level of funding and their presumed sense of responsibility to the public good. When it comes to voting, some people prefer to put their trust in local authorities, such as county supervisors, but most

people recognize that federal organizations have more resources than most local organizations. Thus the federal government has established organizations, including the Federal Elections Commission (FEC)[1] and the Elections Assistance Commission (EAC),[2] to ensure election integrity.

Not everyone trusts government entities, however. For some people, the involvement of government actually increases their distrust. To add to the problem, government agencies—especially those involved with any form of security—tend to keep their methods and operations confidential, which can exacerbate mistrust. Experience has shown that the more transparent and careful an organization is—whether governmental or not—the more likely people are to have confidence in its results.

But who can reassure us that our trust in an external organization is well placed? We often trust organizations because they are bound by law, such as laws against fraud. Then who can we trust to enforce those laws? The law enforcement agencies and the courts. Why should we trust them? And so on.

Eventually we reach a point at which we must place our trust in some entity to provide a reasonable presumption of accuracy, assurance, transparency, and correctness and hope that those presumptions are well founded. This is approximately the same problem as putting our trust in a computing system. At some point we must trust components without formal assurances (Thompson, 1984).

Concurrently, we must balance assurances against risks. What are we likely to lose if our trust is misplaced? I am more willing to trust someone's recommendation if it is for an item that has little consequence to me if it fails. If someone guarantees that I will get my money back or receive a replacement item if something goes wrong, I may be more likely to trust a product based on otherwise weak assurances.

In elections, however, the stakes are often extremely high, so I am less likely to extend my trust. There is no simple recourse for elections. We cannot "get an election back." We do not get a "do over." The potential loss is significant, although not everyone realizes the extent of the risks. The stakes in elections matter at almost every level.

Transparency to Individuals
Very few people know what to look for in assessing the security of a voting system or understand how to look for signs of covert tampering. Very, very few people understand how to examine source codes to determine if something has been embedded or has previously been embedded and has now deleted itself. We know from experience with computer viruses that software can do damage and then delete itself so it is no longer visible (Thompson, 1984). These viruses are extremely difficult to detect.

Most vendors of voting machines do not expose their system implementation or designs to outside scrutiny. Instead, they hide them behind trade-secret protections and become belligerent if anyone tries to examine them. This makes it difficult, if not impossible, for an informed individual who understands the risks to gain self-assurance by looking at the code.

The procedures used in voting are often not fully transparent. They may be described, but they are not transparent in the sense that individual voters can verify every step. For example, in many jurisdictions voting machines are taken home overnight by poll workers and thus are not under continual observation by election judges. These machines may be sealed shut with tamper-*evident* tape, but tamper-evident tape is not completely tamper proof. Furthermore, systems that have screws or other common fasteners that can be removed are not tamper *proof*; nor are systems that have any method of altering the software. Thus in some jurisdictions the procedures are not fully transparent and cannot be audited by individual voters.

People's experience may also lead them to false conclusions. Human beings have a tendency to believe that if something has not happened yet or has not been observed, then it is less likely to occur in the future. For example, we may reason: "I have not died yet, so maybe I am never going to die." Clearly this is fallacious, but generalizing prior personal experience often leads to false conclusions. Not understanding how to balance risk against prior behavior often leads people to act in risky ways, such as smoking, gambling, or driving without wearing a seat belt.

False conclusions can lead to an intuitively appealing inertia with regard to voting systems. People may believe that, because no incidents of catastrophic failure or fraud have been discovered or publicized, they are less likely to occur in the future. That is not the case, however. Not only is failure, such as a type of event not based on prior experience, likely, but it is also possible that fraud may have occurred but was

never noticed (or proven).

Can We Trust Vendors?
If we cannot depend on our own evaluations as individual voters, can we trust the vendors? Experience has shown that trusting vendors is not a good idea. They not only try to hide their processes and their code so that it cannot be independently checked, but also have a vested interest in portraying their systems as infallible. We have no sound basis for trusting them to assure us about their products.

Local election authorities have placed great emphasis on third parties, such as independent laboratories. However, those laboratories have historically been funded by voting machine vendors and thus may have conflicts of interest. So we cannot be sure that they are truly independent.

In addition, testing labs that have been used by voting machine vendors do not publish their procedures, disclose their test sets, or otherwise inform the public about the rigor of their tests. Their prior results have not been published, so the public cannot know when they tested the machines, whether or not the machines passed, what the tests were, and what the results were. Those details are shared only with the voting machine vendors who pay for the testing of their equipment.

Based on what we know, tests by these labs appear to be woefully incomplete. For example, the underlying operating system is assumed to be correct and is not tested. Thus, machines that run voting code on a Windows operating system have a significant gap in system assessment where it is commonly known that vulnerabilities (discovered and yet to be discovered) exist.

We have no sound basis for trusting vendors' assurances about their products.

Many companies that manufacture voting machines hire questionable employees and use questionable practices (Drew, 2007), and faults have been found in certified systems (Wagner et al., 2006). In at least one case, a vendor hired a known felon to work on its products. Some vendors are also alleged to have hired individuals who are not U.S. citizens and may not be fully trust-worthy or may be easily influenced.

The de facto certification framework for the software security industry is the Common Criteria,[3] although its efficacy and value are open to debate. To my knowledge, none of the voting machine vendors has had any of their products certified under the Common Criteria, and certainly not to one of the higher assurance levels. Nor has any of them been certified under ISO standards for security practices or for any other good practices established for developing software. Thus any comprehensive test regime must start with basic examinations, because no standard of good practice can be assumed.

Verification of Voting Systems
To maximize voter assurance, we must consider several factors. First, we must carefully specify what the equipment should do and the procedures for its use. It is hard to believe that this is not done routinely, but it is not. In addition, we should use best methods in the development of voting systems. Many "best practices" have been established by industry, such as the Capability Maturity Model,[4] the ISO standards #90005 and 17799, and the previously mentioned Common Criteria, any of which could be used by voting machine vendors. We should also have open review of the code and testing of voting machines so that external reviewers can examine them for shortfalls and flaws.

We should have open review of the code and testing of voting machines.

Voter Verification
Voter verification is a simple, yet reasonably new concept. The familiar voting experience is to indicate

and cast a ballot, then leave the polling place confident that the vote is counted. However, with new voting technology, such as electronic touch screens, we need voter-verified mechanisms that the majority of voters can use. By verifying that a machine has correctly recorded the vote as cast, we gain confidence in the overall results. Independent verification by the electorate, combined with other checks and safeguards already in place, would make everyone feel more confident that the system is working correctly.

It is not necessary that every single voter verify his or her vote. It is only necessary that a random sample of voters verify their votes to ensure that the machines behave correctly and that errors or fraud are detected on a machine-by-machine basis. Having a simple enough verification mechanism in place that any voter can understand and use makes every voter an election judge and a tester of system assurance. It also helps ensure that the voters who verify their ballots have not been pre-selected. This is one way of randomizing verification.

One very simple voter-verification method is using an optical scanning machine. This scheme, known as a voter-verified paper audit trail (VVPAT), can be accomplished in several ways. For example, a computerized ballot-marking system can be used to create the initial ballot. This system presents the choices in a variety of user-accessible languages or formats, and the interface helps prevent errors, such as a voter failing to cast a vote in each race. It can also present alternate interfaces to accommodate users with disabilities.

Upon completion of voting, the system prints a ballot for optical scanning. Once the voter has verified that the scan matches the ballot, he or she inserts the marked ballot into a scanning machine at the precinct to record the vote. With this system, the user can verify his or her vote and still enjoy the convenience of computerized voting when marking the ballot and tallying all of the votes in the precinct.

To double check that the tally is correct, a random audit must be performed after the polls are closed. This involves performing a manual recount of paper ballots from a set of randomly selected machines or precincts. These counts are then compared against the electronic versions to ensure that there are no significant discrepancies. This type of audit should be done as a matter of course under the careful oversight of the EAC and state and local boards.

The paper and electronic combination can be used with different technologies that have different failure modes. Tampering with one form of ballot cannot easily be replicated in the other, especially with election judges present and random audits and voter self-verification. The paper component is easily understandable by almost every voter, and the system can be implemented with great transparency.

Software Security
We do not yet know how to solve the problem of producing secure, 100 percent correct software for use in voting machines. Figures collected by the US-CERT show that more than 8,000 security vulnerabilities alone were reported in commonly used software in 2006, a time when vendors were trying diligently to keep from introducing new problems into their products. Flaws of any kind are a danger in voting. Votes that are lost or altered by an unexpected failure are as lost as votes corrupted by fraud, and flaws in general are more common than security-specific problems.

In the last few years, Microsoft alone has reportedly spent tens of millions of dollars to retrain its personnel and build new tools to try to eliminate flaws in its software. Despite these efforts, we continue to see reports of new, serious flaws in Microsoft products. Why should we expect voting machine software to be different? In fact, it is not. Unfortunately, it is not even produced with the same care as some software games and word processing software.

Moreover, there are theoretical limits to software assurance as a process. Many believe it is easy, or at least straightforward, to examine software systems to detect flaws. Quite the contrary! Those of us who work in security and software testing know that there is no way to guarantee that software does not have hidden flaws. And, as government agencies use more and more software produced outside the United States, this is becoming a major national security issue, not only related to voting equipment.

Despite many, many decades of effort and funding, our national agencies and research labs have not been able to develop systems that can detect whether or not there is hidden functionality in selected code. And despite decades of effort by industry and academia, we still do not know how to produce large, error-free software artifacts. Thus it seems likely that our shortcomings in assurance methodologies will continue to impose limits on how well we can assure software in voting systems.

Clearly, voting systems are far too important to operate on faulty software, or even on systems that are not designed for high assurance. Because of the necessarily distributed nature of election management, we must find a way to provide local election officials with credible information about voting system security before they purchase and operate new systems. Recent efforts by the EAC (and FEC before them) to implement testing and certification procedures for voting systems have been well intentioned, but so far ineffective. In the meantime, VVPAT voting methods would go a long way toward restoring voter confidence.

Conclusion

It is almost inevitable that some people will be unhappy and believe that an election went the wrong way. Usually these are supporters of the losers, but some results may also make winners feel uncomfortable. Missing votes, user complaints of disenfranchisement, and unexpectedly lopsided vote totals are inappropriate no matter which party or candidate benefits. Whether these anomalies are the results of tampering or unex-pected errors, they add to voters' disillusionment. The voters in this country (and arguably, every country) deserve better. Moreover, our form of government demands better, because our ability to elect our representatives confidently is the foundation of our republic.

The changes triggered by the 2000 election and signaled by HAVA will not happen overnight, and even those changes will not completely solve all of the problems. However, voting system assurance is possible through the rigorous application of engineering practices in the development, certification, and operation of voting machines, coupled with the appropriate use of audits and operational care. We must remain involved and diligent. It is our voting system, and it will only be trustworthy if we demand that it is.

Acknowledgment

The author thanks Alec Yasinsac for his invaluable assistance in preparing this paper for publication.

References

Drew, C. 2007. U.S. bars lab from testing electronic voting. New York Times, January 4, 2007.

Elias, T. 2007. Is it high time for an e-voting test in California? The Daily Independent, Friday, January 26, 2007.

Greene, S. 2006. "With Voting Machine Company Now Bankrupt, CEO Speaks Out: No Vendor Has a System That Voters Can Trust!" electionline.org, April 5, 2006.

HAVA (Help America Vote Act). 2002. Public Law 107-252, 107th Congress, United States.

Kohut, A. 2006. Public Concern about the Vote Count and Uncertainty about Electronic Voting Machines. Washington, D.C.: Pew Research Center for the People & the Press, November 2006.

Thompson, K. 1984. Reflections on trusting trust. Communications of the ACM 27(8): 761–763, Aug. 1984. Also in ACM Turing Award Lectures: The First Twenty Years 1965–1985, Copyright 1987; and Computers Under Attack: Intruders, Worms, and Viruses, Copyright 1990, both by ACM Press. Available online at http://www.acm.org/classics/sep95/.

UMd (University of Maryland). 2002. An Evaluation of Maryland's New Voting Machines. College Park, Md.: Human-Computer Interaction Lab, Center for American Politics and Citizenship, University of Maryland.

Wagner, D., M. Bishop, D. Jefferson, C. Karlof, and N. Sastry. 2006. Security Analysis of the Diebold AccuBasic Interpreter. Berkeley, Calif.: University of California Press.

FOOTNOTES

1 http://www.fec.gov/.
2 http://www.eac.gov.
3 Information available at http://www.commoncriteriaportal.org/ and elsewhere on the Internet.
4 http://www.sei.cmu.edu/cmmi/.
5 http://www.iso.org/iso/en/iso9000-14000/index.html.

About the Author

Eugene H. Spafford is executive director of the Center for Education and Research in Information Assurance and Security and a professor of computer sciences at Purdue University.