RESPONSES TO QUESTIONS FOR THE RECORD
DAVID WAGNER, PH.D.[1]
COMPUTER SCIENCE DIVISION
UNIVERSITY OF CALIFORNIA, BERKELEY
SEPTEMBER 11, 2006

Responses to "Questions for the Record Submitted by Chairman Ehlers and Chairman Boehlert to Dr. David Wagner"

1. How do you think the sections of the 2005 Voluntary Voting Systems Guidelines (VVSG) that deal with security should be improved?

I recommend sweeping changes to how the 2005 Voluntary Voting Systems Guidelines (VVSG) deal with security, to bring them up to date with fundamental changes over the past decade in how voting systems are built. The 2007 VVSG are in the process of being drafted, and I propose several suggestions for consideration.

- *Require that systems provide voter-verified paper records.* The single most effective step that the VVSG could take to improve security would be to stop certifying new voting systems that do not provide a voter-verified paper record. The VVSG could also be revised to require that the use procedures provided by the vendor specify how to perform a routine manual audit of these paper records.

  Given the current state of the art, there is no known way to provide a comparable level of security without voter-verified paper records. In the long run, as technology advances, it may be possible to develop alternative voting technologies that provide an equal or greater level of security without using paper. Consequently, it may be appropriate to structure the VVSG to permit other systems that demonstrably provide an equal or greater level of security as voter-verified paper records with manual audits. However, any such provision would need to be accompanied by a new process for determining which systems meet this criteria. The current evaluation and testing process is not capable of making these determinations with any credibility; major reforms of the current processes would be required before such a provision would be safe to add. Adding such a provision without accompanying reform of the process used to evaluate which systems qualify for the exception would eliminate much of the benefit of a requirement for voter-verified paper records. In addition, it should be expected that evaluating the security of systems that do not use voter-verified paper records will be considerably more expensive and difficult than evaluating systems that use voter-verified paper records, due to the fact that paperless systems do not record a permanent copy of the voter's intent that the voter can verify.

- *Begin enforcing existing requirements.* At present, many of the security requirements in the 2005 VVSG are not enforced or tested by the federal qualification process. While the existing requirements of the VVSG are, for the most part, a fairly reasonable start at specifying security requirements for a voting system, the lack of enforcement renders these well-intentioned requirements ineffective.

[1] This work was supported by the National Science Foundation under Grant No. CNS-052431 (ACCURATE). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation. I do not speak for UC Berkeley, the National Science Foundation, or any other organization. Affiliations are provided for identification purposes only.

The VVSG do not specify any specific testing procedure for many of the security requirements, and perhaps as a consequence, the federal testing labs apparently do not perform an independent analysis of whether these requirements are met. Instead, the testing labs seem to concentrate their efforts on requirements for which there is a concrete testing procedure defined in the VVSG. We now know of multiple examples where the federal testing labs have approved voting systems that contain violations of the VVSG[1].

- *Create faster ways to investigate and act on experience from the field.* At present, the EAC has no way to respond quickly to new discoveries about the security of deployed voting systems. Currently, the only mechanism the EAC has to affect the machines that voters vote on is to revise the VVSG. However, these revisions take an extremely long time to take effect. For instance, the next revision of the VVSG is not scheduled until 2007. Moreover, the 2007 VVSG are not expected to take effect until 2009. Furthermore, when the 2007 VVSG do go into effect in 2009, they will only affect newly developed or modified systems submitted for certification after that date. Any systems that had been already certified or already deployed at that time would be grandfathered. Consequently, any new provisions in the 2007 VVSG will only affect systems purchased after 2009, and possibly only systems that were both developed and purchased after 2009. Because jurisdictions purchase new systems only rarely—perhaps once a decade or so, at best—any revisions to the VVSG that the EAC wished to make today might not have any impact on the machines that a majority of Americans vote on until 2015 or so.

  Moreover, the EAC has no formalized, systematic way to gather data from the field about the performance of voting systems or to track incidents and failures across the country.

  In comparison, the aviation industry has more effective mechanisms for investigating and responding to new discoveries about threats to aviation safety. Whenever a plane crash or other serious in-flight anomaly occurs, federal investigators immediately investigate the cause of the failure. If serious problems are found, federal regulators have the authority to require that corrective action be taken immediately, if necessary. The consequence is that federal authorities have the ability to respond to serious problems that affect aviation safety in a matter of months. The EAC lacks any corresponding capability to investigate or respond to voting system failures.

  It would help to create ways to investigate voting system failures, to require reporting of election incidents, to gather data from the field and quantitatively measure the rate of failures, to update voting standards more frequently in response to this data, and to require timely adherence to the standards[2].

  Also, it would help to establish a process to decertify voting systems that are certified and then are subsequently discovered to have security flaws or to violate the standards. It would help if the EAC were to exercise its authority to decertify systems when they are found to have security vulnerabilities.

- *Require some additional safeguards recommended by security experts.* Many security experts have recommended several additional safeguards: banning wireless communications in voting systems; banning some forms of interpreted code; banning code stored on removable storage media. These would not on their own fix all the security problems we are currently experiencing, but they would help address some known gaps in the standards.

Do you think that the way in which security for voting systems is tested needs to change? If so, how, and if not, why not?

Yes. The current process is not working: systems with serious security vulnerabilities are getting approved. I suggest several reforms.

- *Convene a panel of security experts to conduct independent security evaluations of every system submitted for certification.* Each time a voting system is submitted to the federal qualification process, the EAC could convene a panel of leading security experts from both academia and industry to perform an independent security analysis of the system. Independent security evaluations are standard practice in the field of computer security; the election industry has lagged behind the rest of the field in this respect.

  Over the past few years, external experts have been much more effective at finding security flaws and assessing the security of today's e-voting systems than the federal testing labs. Consequently, it makes sense to enlist those who have demonstrated skill at finding security vulnerabilities in voting systems, so that we know about the flaws and can take appropriate action before the systems are deployed in the field. For instance, in 2003 four academics found more security flaws in one voting system in 48 hours of examination of the voting software than the federal testing labs had in the years that the system was deployed. In 2005, a Finnish security researcher found two significant security vulnerabilities after approximately one week of study of a voting system, upon the request of a county election official in Florida. In 2006, the same Finnish researcher found another serious security vulnerability after another week of study of the same voting system, at the request of a county election official in Utah. Independent security evaluations could help reduce the chances of approving and deploying a flawed system.

  Given that many have lost faith in the ability of federal testing labs to evaluate the security of voting systems, independent security evaluations would provide an independent check on the federal testing labs. Because the effectiveness of an independent security evaluation is highly dependent upon the skills of the participants, it is important that panelists be chosen from among the best minds in computer security. To this end, I would recommend that that the EAC consult with the ACCURATE project to identify potential panelists. The panel should have full access to all technical information about the voting system, including all source code. The panel should also have full access to a working unit of the voting system, and the authority and ability to physically inspect and run tests on that unit. The panel should be asked to write a report of their findings, and the report should be made public in its entirety. If necessary, the vendor's proprietary interests can be protected, while preserving transparency and the independence of the evaluators, through an appropriate nondisclosure agreement.

- *Require vendors to disclose the source code of all voting system software by a specified future date.* The use of secret software has contributed to a loss of transparency and eliminated opportunities for public oversight of important parts of the machinery of our elections[3]. This secretiveness has contributed to a loss of confidence in the voting systems. The best way to remedy this would be to require that vendors make all source code, and other technical information about the design and construction of their voting machines, publicly available for all interested parties to examine[4]. Vendors would still enjoy the protection of patent and copyright law but would be required to forfeit trade secrecy in their software to field systems in federal elections.

Some transition strategy may needed to phase in this requirement. One possibility is to specify a date several years in the future after which source code to voting systems would be required to be disclosed and provide advance notice to vendors of that date. In the short term, source code might be required to be disclosed to any accredited security expert who is willing to sign appropriate nondisclosure agreements.

- *Eliminate the COTS loophole.* The standards currently contain an exception that exempts commercial off-the-shelf software (COTS) from some of the testing. Because COTS software has been implicated in some recent security vulnerabilities, I believe there is a good argument for eliminating this exception.

- *Eliminate conflicts of interest; ensure that evaluators are truly independent.* At present, the federal testing labs work for the vendors: they are paid and selected by the voting vendors. We need some other mechanism that better ensures the independence of the testing labs.

  One possibility would be for the testing labs to be paid by the federal government, with vendors required to reimburse the government for all costs incurred. For instance, in California the state has set up an escrow account for each vendor. The vendor is required to deposit sufficient funds to cover all the costs of certification testing into this account; when the state hires consultants or other experts, they are paid out of this escrow account. The federal government could use a similar system. This would make it clear that labs work for the federal government and have a fiduciary responsibility to the citizenry, not to the vendor.

  It may be possible to devise creative new approaches that rely on market forces to make testing more effective. For instance, if federal labs had to pay damages when a voting system they approved turned out to be insecure, they would have an incentive to make their testing processes as effective as possible. One possibility might be to require federal labs to carry insurance and give all citizens standing to sue the labs for approving insecure voting systems, setting the damages for endangering democracy at a high dollar amount. Federal approval of a voting system might mean far more if testing labs needed to keep their insurance premiums down in order to remain profitable. It is not clear whether such an approach can be made workable, but new incentive structures may be worth exploring.

- *Make all reports from the testing labs public.* Today, the results from the federal testing labs are not made available to the public. The labs consider them proprietary and the property of the vendor. If a system fails to gain the testing lab's approval, this fact is not disclosed to anyone other than the vendor who paid for the testing.

  I recommend that the results of all testing at the federal level be disclosed to the public. All reports produced by the testing labs should be published in full, whether the systems pass or fail.

- *Enforce all security requirements in the standards.* As mentioned earlier, many security requirements are never tested and consequently are not enforced. Security evaluation of voting systems should change so that all security requirements are assessed. We should expect and require testing labs to fail any voting system if they cannot demonstrate that it meets all security requirements.

2. Is computer security testing different from other types of conformance testing, and if so, how? Has this type of testing ever been performed on voting equipment and if so, what were the results? Should this type of testing be performed routinely on voting equipment?

Yes, security evaluation is different from other types of conformance testing. Conformance testing—commonly also known under the name "functionality testing" or "black-box testing"—is concerned with ensuring that the system will respond in certain ways under ordinary operating conditions. This makes conformance testing fairly straightforward: the best simulates ordinary operating conditions and then checks that the system responds as desired under these conditions. For instance, if we want to test that a voting system correctly counts write-in votes under normal operating conditions, then we can run a mock election, cast several write-in votes, and confirm that they are counted correctly. As this example illustrates, conformance testing is often fairly straightforward.

In contrast, security evaluation is concerned with ensuring that the system will not misbehave when it is intentionally misused. Thus, ordinary conformance testing is concerned with how the system behaves under normal conditions, while security evaluation is concerned with how it behaves under abnormal conditions. Unfortunately, it is very difficult to predict how an attacker might try to misuse the system. If we could predict how the attacker were going to misuse the system, then we could simulate such misuse and observe whether the system is able to respond appropriately. However, usually we do not know how an attacker might try to misuse the system, and there are too many ways that an attacker might try to misuse the system to exhaustively enumerate them all. Consequently, there is no way to simulate how the system reacts to these kinds of unanticipated attacks. This makes security evaluation more difficult than ordinary standard conformance testing.

For these reasons, standard conformance testing practices are not effective at evaluating whether a system is secure or not. Security practitioners are familiar with this phenomenom[5]. As a result, when experienced practitioners need to evaluate the security of some software, they normally use discipline-specific methods chosen to be effective for security purposes, instead of just relying on testing. These methods always include some form of adversarial analysis, which may include elements of threat assessment, source code review, architectural review, penetration analysis, and red teaming. Security practitioners also understand that, to be most effective, adversarial analysis should be performed by security experts who are neutral and independent. This process of adversarial analysis, when performed by independent security experts, is sometimes known under the name "independent security evaluation". Use of these adversarial analysis methods is routine practice in industries where security is mission-critical.

Yes, these security evaluation practices have been applied, on a limited basis, to several voting systems. In each case, serious security flaws were found.

- In 2003, researchers from Johns Hopkins and Rice Universities undertook an adversarial analysis and source code review of voting software used in Diebold touchscreen voting machines[6]. They found numerous security vulnerabilities.

- In 2004, a security consulting company (RABA Technologies) performed an independent security evaluation of Diebold voting systems and found several security vulnerabilities[7].

- In 2005, Finnish researcher Harri Hursti applied source code analysis and testing to discover and confirm two security vulnerabilities in an optical scan machine manufactured by Diebold[8].

- In 2006, I and several other security experts analyzed source code provided by Diebold as part of our independent security evaluation of Diebold systems[9]. We confirmed that Hursti's vulnerabilities were present in both Diebold optical scan and touchscreen machines. We also found 16 other security defects that had not been previously known.

- In 2006, Hursti was asked to examine a Diebold touchscreen machine, and he discovered another serious security vulnerability using adversarial analysis[10].

In each case, the use of practices specific to the field of computer security was central to the effectiveness of these security evaluations. As far as I can tell, none of these security vulnerabilities had been previously discovered by the federal testing labs, perhaps because the labs were focused on standard conformance testing and failed to use methods more appropriate to security evaluation[11].

Yes, these security-specific evaluation methods should be applied routinely to voting systems. They are the best tools we have for weeding out insecure voting systems, for proactively finding and fixing security vulnerabilities in voting systems before they are deployed, and for increasing confidence in the security of these systems.

It is worth mentioning that the term "testing" has a more specific meaning in the computer science jargon than its everyday meaning. Someone who is not a computer specialist might use the word "testing" to describe any method for evaluating the quality of software or for finding software defects. In contrast, computer scientists use the term "testing" more narrowly to refer to one specific method for evaluating software quality: among computer scientists, the unqualified term "test" is often viewed as a synonym for "black-box testing", "functionality testing", or "conformance testing". Computer scientists would say that "testing" is just one method of assessing the quality of software, but that there are others, as well. When it comes to security, those other methods are usually more effective than "testing". Because of the potential for confusion, I will avoid use of the unqualified word "testing"; I will use terms like "functionality testing" to refer to one specific method of evaluating software quality, and terms like "evaluation" to refer to the broad goal of evaluating software quality and finding software defects.

3. In your written testimony, you stated that functionality testing is not as good as discipline-specific testing. Please explain the difference between functionality and discipline-specific testing, and why you believe discipline-specific testing should be used for voting equipment.

"Functionality testing" is a synonym for "black-box testing" or "conformance testing". Thus, my response to Question 2 is relevant to this question as well.

As I mentioned, security practitioners have developed discipline-specific methods—methods that are suited to the discipline of computer security—for evaluating the security of computer systems. These include source code analysis, independent security analysis, architecture and design reviews, and red teaming. Functionality testing verifies that a machine does what it is supposed to do, when it isn't under attack; in contrast, these security evaluation methods verify that a machine does not do what it isn't supposed to do, even when it is under attack. These discipline-specific methods should be used on voting equipment in addition to functionality testing, because they are the best known way to assess the security of such systems.

The discipline of usability has also developed its own discipline-specific methods for evaluating the usability and accessibility of computer systems, including user testing with actual voters and pollworkers as well as heuristic evaluation by usability and accessibility experts. These methods specifically cater to human factors concerns and are designed to evaluate how the software influences interactions between humans and computers. These methods are focused less on functional requirements (e.g., can the system display candidate names in a bold font?) and more on assessing performance via quantitative metrics of usability. These discipline-specific methods should be used for voting equipment, because they are the best known way to assess the usability and accessibility of such systems.

4. Mr. Groh and Ms. Lamone expressed concerns about the use of the voter-verifiable paper audit trail. These concerns included the additional costs to jurisdictions of implementing these systems, and the accessibility of such technologies to the disabled community. Ms. Lamone also cited a Maryland study that indicated that the paper trail, in addition to other verification technologies, was not ready for primetime. Do you agree with these concerns? If so, why, and if not, why not?

In short: I agree with the concerns about cost; I do not agree with the concerns about accessibility; I do not agree with Ms. Lamone's characterization of the Maryland study. I provide my reasoning below.

- I do share Mr. Groh and Ms. Lamone's concerns about the costs of implementing systems that support voter-verified paper records. Approximately 15 states have purchased paperless voting systems that do not provide voter-verified paper records[12]. Some of these paperless voting systems can be retrofitted to produce a voter-verified paper trail, but in some cases these systems cannot be easily upgraded or retrofitted with a paper trail. Even when it is possible, retrofitting is not cheap. Replacement is even more expensive, as it involves throwing away equipment and replacing it with more modern equipment. It is certainly understandable why states who have made a significant investment into a particular voting system would be reluctant to scrap these systems and incur significant costs in replacing them. It is unfortunate that some states bought paperless voting systems without realizing the security, reliability, and transparency consequences of that action.

  The costs would vary widely from state to state. Currently, 27 states require by law that all voting systems produce voter-verified paper records, and another 8 states have deployed voting systems with voter-verified paper records even though state law does not require it. In total, 35 states (70% of states) have voting systems that already produce a paper audit trail and would not need to be upgraded or replaced. Those 35 states would not incur any cost. The remaining 15 states (30%) do not consistently use systems with a paper audit trail statewide. In those states, some or all of the voting equipment in the polling places would need to be upgraded, retrofitted, or replaced. On the other hand, equipment used for scanning absentee (mail-in) ballots, which account for 30-40% of the vote in many states, would not need to be changed.

  Even within this class of 15 states, costs would vary by state. At one extreme, some states use paperless DREs throughout the state, and all of those DREs in every county would need to be upgraded, retrofitted, or replaced. As best as I can tell, there appear to be 5 states (DE, GA, LA, MD, SC) in this category. Of those 5 states, 2 (GA, MD) use DREs that would need to be completely replaced, because there is no good way to upgrade or retrofit them with a paper trail; 2 (LA, SC) use DREs for which an approved printer add-on is already on the market; and I do not know whether retrofitting is possible in the remaining state (DE). Obviously, replacing all DREs is the most expensive possible case. At the other extreme, in some states the voting equipment is not uniform throughout the state and costs would be less in some counties than in others. For instance, approximately 52 of 67 Florida's counties use optical scan voting machines plus one accessible voting system (DRE or ballot marking device) per polling place; upgrades for those counties would be less expensive, because the optical scan machines would not need to be upgraded, retrofitted, or replaced.

Costs will also vary according to the system that is in use. Many modern DREs (e.g., the Diebold TSx, ES&S iVotronic, Sequoia Edge, and Hart-Intercivic eSlate) can be upgraded to produce a paper trail: approved printer units are available on the market. Upgrading these DREs to add a printer might cost approximately $500-$2000 per DRE, depending on the vendor. Some older DREs (e.g., the Diebold TS) cannot easily be upgraded or retrofitted with a paper trail, and would have to be replaced with all new equipment. Buying new DREs normally costs about $3000-$5000 per DRE. However, in some cases it may be cheaper to replace the paperless DREs with a hybrid system using optically scanned paper ballots. These hybrid systems require purchasing one optical scan machine plus one accessible voting machine (DRE with VVPAT or ballot marking device) per precinct, and this equipment typically costs in the ballpark of $10,000-$12,000 per precinct. Because an all-DRE solution usually requires several DREs per precinct, hybrid systems using optical scanners may come out cheaper. The cost advantages of hybrid systems are more pronounced in states that require DREs to display a full-face ballot, because full-faced DREs are significantly more expensive than standard DREs[13]. I would encourage jurisdictions to consider all available options.

In summary, I do not know what the total costs might be, but I share Mr. Groh and Ms. Lamone's concerns that the costs of implementing a voter-verified paper trail will be significant in some states.

- I do not agree with their concerns about the accessibility of these voting systems to the disabled community. The disabled community has praised the development of touchscreen voting systems as providing major improvements in accessibility, and rightly so: the accessibility benefits are significant and real. However, voter-verified paper records are in no way incompatible with these benefits. Today, every major vendor who offers a touchscreen voting machine also offers a version of that touchscreen machine that produces a voter-verified paper record. Those VVPAT-enabled versions provide the same accessibility support—audio interfaces, high-contrast displays, sip-and-puff devices, booths designed for wheelchair voters, and so on—as their paperless brethren do. Adding a printer makes the machine no less accessible.

  I believe security and accessibility do not need to be in conflict; I believe we can have both. This is fortunate, because I believe both security and accessibility are important goals.

  I understand that one concern is that visually impaired voters will not be able to independently verify what is printed on the voter-verified paper record. This concern is valid, but I do not consider it a persuasive argument against voter-verified paper records. If a blind voter does not trust the voting machine to work correctly, then it is true that they have no way to independently verify that their vote has been recorded correctly. In other words, blind voters must rely upon the voting software to work correctly, and they are vulnerable to software failures; they have no independent means of checking that the software is working correctly. This situation is truly unfortunate. However, this is the case for all currently available voting technologies, whether they print a paper record or not. If the machine prints nothing, then the blind voter still cannot independently verify that their vote has been recorded correctly on electronic storage. To put it another way, with paperless voting machines, neither sighted voters nor blind voters have any chance to independently verify their vote; with voter-verified paper records, sighted voters can independently verify their vote, but blind voters cannot. Voter-verified paper records do not make the independent verification problem any worse for blind voters; they just fail to make things better.

The policy question is whether it is valuable to improve security and reliability for most voters, even if there are some voters who are not helped by these measures (but are not harmed by them, either) and remain without any means of independent verification.

- I do not agree with Ms. Lamone's characterization of the Maryland study. At present, Maryland uses a paperless touchscreen voting machine, called the Diebold TS. The Maryland study was commissioned to study whether there exists any technology currently on the market that could be used to upgrade or retrofit the Diebold TS with a way for voters to independently verify that their vote was recorded, and to evaluate whether any of these are ready for use in real elections. The Maryland study was specifically limited to studying methods of upgrading or retrofitting the Diebold TS; replacement was out of scope for the study. The conclusion of the study was that there was no good way of upgrading the Diebold TS that would be ready for use in the near future. I have read the study carefully and I agree with that conclusion. I agree with Ms. Lamone that the study was "very thorough" and "provided some very valuable information."

  However, I disagree with Ms. Lamone's characterization of the study as finding that "the paper trail" was not "ready for primetime." In fact, the Maryland study's findings were more narrow than that. The Maryland study was asked not to consider any technology that would require replacing Maryland's Diebold TS machines; they were asked to consider only technology for upgrading those machines, and they did so. It is indeed justified to conclude from the study that none of the systems for upgrading the Diebold TS are "ready for primetime." However, the study says nothing about the viability of other, more modern voting systems that do provide a voter-verified paper trail. The correct conclusion to draw from the Maryland study is that if Maryland wants to adopt voter-verified paper records, they will need to replace their existing Diebold TS machines; retrofitting is not a viable option. The study says nothing about whether existing, deployed systems that provide a paper trail are ready for primetime. I believe there are existing paper-trail systems that are already ready for primetime.

  Maryland is in an admittedly difficult position. Maryland was one of the first states to adopt touchscreen voting systems, and while the Diebold TS machines they bought were thought by some to be adequate at the time, at present the Diebold TS machines are no longer the most current technology. The Diebold TS was not designed to provide a paper trail. Its successor, the Diebold TSx, does provide a voter-verified paper audit trail. The other major voting system vendors also sell voting machines that do provide a paper trail. Not all states are in the same position that Maryland is in: many states already use systems with a voter-verified paper trail; and some states have voting systems that do not currently provide a voter-verified paper trail, but that can be upgraded or retrofitted to provide a paper trail.

5. The 2005 VVSG contains an appendix on independent dual verification systems that could perform the same functions as a voter-verifiable paper audit trail. Is this technology being used in voting systems today or is more research needed to make it operational? What are the advantages and disadvantages of this technology? To what extent are there other technologies that could perform the same function as a voter-verifiable paper audit trail?

No, this technology is not being used today in any deployed voting system that I am aware of. More research would be needed to determine whether the approach can be made operational. The future of this approach is uncertain at this point.

The advantages and disadvantages of any particular system will depend on how that system is designed and implemented. It is difficult to comment on advantages and disadvantages in the absence of a fully implemented system. I can only speculate.

One potential disadvantage is that evaluating whether these systems meet the security requirements is likely to be significantly more expensive for paperless independent dual verification systems than for systems producing a voter-verified paper record, both because the certification process would need to be overhauled, and because assessing whether paperless independent dual verification systems are secure is inherently more difficult than assessing whether systems with a paper trail meet their security goals. Another potential disadvantage of paperless independent dual verification systems is that it may be harder for voters who do not have a degree in computer science to know whether they should trust those systems. One motivation for seeking paperless systems is that eliminating the need to handle or store paper could make election administration more efficient. Also, ideally such a system might provide visually impaired voters with a way to independently verify their vote, which would be a significant advantage. Unfortunately, no such method is known at present.

At present, it is an open question whether it will be possible to develop a paperless voting system that can perform the same function as a voter-verified paper trail. There does not appear to be any firm consensus among computer scientists on whether such an alternative is even possible, given the current state of technology; on what directions are most promising to explore; or on how far off this goal may be. I believe that more research is warranted, but that we should not expect deployable replacements for paper any time soon.

6. Have you conducted any studies of the problems/deficiencies of paper based systems?

Yes. I have conducted studies that revealed some problems and deficiencies in certain paper-based systems. I have not attempted to undertake any study to exhaustively categorize all possible problems or deficiencies that can arise with paper-based systems. Of course, the history of paper-based elections in this country dates back at least two hundred years, and it is well-known that they can be susceptible to certain kinds of problems (e.g., problems in the handling, transportation, or storage of paper ballots) if elections are not well-administered.

Is your support for a voter verified paper record principally motivated by confidence in paper based systems or a lack of confidence in direct recording electronic systems? If the former, what is the source of this confidence? If the latter, on what basis do you conclude that paper based systems are necessarily superior?

My support for voter-verified paper records is motivated both by confidence in paper-based elections (if they are administered well) and by my lack of confidence in paperless DRE machines.

My confidence in systems that produce voter-verified paper records and include routine manual audits is based on my study of these systems and on analysis of their security properties. My confidence in these systems is based on the ability of voters to verify for themselves that their vote was recorded as they intended, and on the ability of observers to verify that votes were counted correctly and to exercise effective oversight of the process.

My lack of confidence in paperless DRE machines is based on my study of these systems, on analysis of these systems in the open literature[14], and on the documented security flaws and failures of these systems. For instance, the Brennan Center report found that with paperless DRE machines, a single malicious individual with insider access may be able to switch votes, perhaps undetected, and potentially swing an election. The analysis in the Brennan Center report also found that systems that produce voter-verified paper records and include routine manual audits are significantly more secure against these threats than paperless DRE machines.

7. Do you foresee any problems that might arise in jurisdictions utilizing a voting system that attaches printers to Direct Record Electronic voting machines? What do you think they might be?

Yes. There are several issues such jurisdictions may want to be aware of.

First, the introduction of printers raises questions of printer jams and the reliability of these devices. California's solution to this problem has been to adopt volume testing, where approximately 10,000 ballots are cast on 50-100 machines in a mock election. Volume testing seems to be effective in weeding out unreliable machines and improving the reliability of voting machines—including their susceptibility to printer jams. The first such volume test found serious printer jam problems in one voting system; fortunately, the vendor was able to correct those problems, and subsequently their system passed the volume testing with no serious problems. California has now certified several DRE voting machines that come with an printer, and these systems appear to provide a satisfactory degree of reliability.

Second, a voter-verified paper record is only effective in proportion to the number of voters who actually verify the paper record as they cast their ballot[15]. Consequently, jurisdictions may wish to consider undertaking voter education to inform voters of the importance of checking the accuracy of the voter-verified paper record.

Third, there is no point in printing a voter-verified paper record if those paper records will never be used or examined by election officials for their intended purpose, i.e., to check vote counts. For this reason, it is important that the jurisdiction create procedures specifying the conditions under which those paper records will be inspected, and what will be done in case of a discrepancy between the paper record and the electronic record. My own recommendation is that jurisdictions adopt routine manual audits; that discrepancies trigger an investigation; that any unexplained discrepancies discovered trigger a manual recount; and that in the event of a discrepancy between the electronic record and paper record, the paper record verified by the voter should have a (rebuttable) presumption of accuracy unless there is some specific reason to believe that the paper records are inaccurate or incomplete.

Fourth, in any election system that uses paper, the handling, transportation, and storage of the paper records is crucial. It is important that jurisdictions establish procedures to establish a good chain of custody for paper ballots and paper trails. For instance, analysis performed by the Brennan Center shows that, if the chain of custody is done poorly, jurisdictions may still be vulnerable to fraud, no matter what voting technology they use.

Finally, and most importantly, the success of an election is determined by more than just technology: it depends crucially on the people who run the election and the processes and procedures they use. Effective and competent election administration is crucial—and printers do not eliminate this important requirement.

Responses to "Questions for the Record Submitted by Democratic Members, Committee on Science"

1. Dr. Wagner, to what extent do voting system security vulnerabilities outlined in the Brennan Center Study reflect weaknesses in the 2002 standards and current certification process? To what extent have those weaknesses been addressed in the 2005 version of the voting systems guidelines and proposed certification process?

The threats outlined in the Brennan Center study reflect significant gaps in the 2002 standards and in the current certification process. The Brennan Center study identified potential threats to voting systems that are not addressed by the 2002 standards or by the current certification process.

Those gaps have not been addressed in the 2005 standards or the certification process it proposes. The Brennan Center study suggested six concrete recommendations to improve the security of elections. None of those are required or recommended by the 2005 standards. In some cases, the 2005 standards takes stances that are directly at odds with the recommendations of the Brennan Center study. For instance, the Brennan Center study recommended banning all wireless communications, yet the 2005 standards explicitly allow wireless communications under certain conditions. One lesson from the Brennan Center study is that the best defense against these threats is the use of voter-verified paper records with routine manual audits; however, the 2005 standards do not require voter-verified paper records or manual audits. If voter-verified paper records are not in place, the Brennan Center recommended that parallel testing be used as a stop-gap; however, the 2005 standards do not require parallel testing, and very few states currently undertake the effort (and expense) of parallel testing.

2. Dr. Wagner, what additional measures need to be taken at the federal level to reduce the incidence of voting system vulnerabilities and problems across the US.

Please see to my answers to Question 1, starting on page 1, for detailed suggestions.

The most significant step that could be taken is to mandate that all voting systems provide voter-verified paper records, and that jurisdictions perform routine manual audits of these records. Also, it would help to conduct more rigorous testing of voting machines, performed by truly independent authorities, using testing methods based on the best scientific and engineering understanding from each applicable discipline and performed by experts from each relevant field; to invite outside security experts to perform independent security evaluations of all voting systems before certification; to increase transparency surrounding the federal testing and qualification process; to begin enforcing the existing security requirements already in the standards; to strengthen the security requirements and testing processes so they reflect the latest understanding of voting systems; and to disclose the source code of all voting systems.

3. Dr. Wagner, why do you believe that electronic voting machines can not be trusted?

If the electronic voting machines are accompanied by a voter-verified paper trail and routine manual audits, and if they are used properly, I believe that they can be trusted. Under these circumstances, they may offer some significant advantages.

However, I do not believe that paperless electronic voting machines can be trusted. The evidence that would be required to trust them is nowhere to be found.

It is beyond the state of the art to verify that the software and hardware used in voting systems will work correctly on election day. For instance, how do we know that a programmer at the vendor has not introduced malicious logic into the voting system? The short answer is that we don't. Malicious logic that has been introduced into a voting system could, for instance, switch 5% of the votes away from one candidate and to the benefit of some other candidate; in a close race, this might make the difference between winning and losing, and such an attack might be very hard to detect. At present, we have no good ways to gain any confidence that our voting systems are free of malicious code; that is beyond the state of the art[16]. Consequently, it seems there is little alternative but to assume that, for all we know, our voting systems could potentially be tampered with to introduce malicious code that will be triggered in some future election.

A second significant concern arises due to the possibility of defects unintentionally introduced into voting systems. Modern electronic voting systems are a highly complex assembly of software and hardware, and there are many things that can go wrong. It is not possible, given the current state of technology, to verify that voting systems are free of defects, flaws, and bugs, or to verify that they will record and count votes correctly on election day; given the complexity of modern voting systems, this is beyond the state of the art.

Consequently, at the moment there seems to be little or no rational basis for confidence in paperless electronic voting machines[17]. In the end, it's not up to voters to take it on faith that the equipment is performing correctly; it's up to vendors and election officials to prove it.

4. Dr. Wagner, why is it that most security experts and computer scientists believe it is necessary to regularly audit voter verified paper trails?

Routine audits are crucial if we are to trust electronic voting[18] [19]. With both DREs and optically scanned paper ballots, it is important to routinely spot-check the paper records against their electronic counterparts. As I explained in my response to Question 3, there is no basis for confidence in the electronic records produced by electronic voting systems—we cannot know, a priori, whether they are correct or not. Given the stakes, we have to be prepared for the worst: that the electronic records may be inaccurate or corrupted. The purpose of a manual audit of the voter-verified paper records is to confirm whether or not the electronic records match the paper records verified by the voter.

The paper records verified by the voter are the only records that we can rely upon to be accurate: they are the only hardcopy record of voter intent, and they are the only records that the voter has the chance to inspect for herself. It would be perfectly adequate, from a security point of view, to simply discard the electronic records and to manually count all of the voter-verified paper records (without the assistance of computers). Such a 100% manual count would produce results that could not be corrupted by computer intrusions, malicious logic, or software defects. However, manual counting of paper records is labor-intensive and costly. Given the number of contests on a typical American ballot today, routine 100% manual counts are probably not economically viable.

To address these concerns, voting experts have devised an alternative that preserves the cost-efficiency of electronic vote counting with the trustworthiness of 100% manual counts[20]. This alternative is based around machines that produce voter-verified paper records along with routine manual audits. During the audit, the paper records from some percentage (perhaps 1% or 5%) of the precincts are manually counted; then the paper tallies are compared to electronic tallies. If they match exactly in all cases, then this provides evidence that the electronic vote-counting software produced the same vote totals that a 100% manual count would have produced, which provides a rational basis for confidence in the election outcome. On the other hand, any mismatches discovered during the audit indicate that something has gone wrong. This provides an opportunity to identify the problem and remedy it, if possible, or to perform a 100% manual recount if the problem cannot be identified.

Consequently, routine manual audits are the best way to ensure that the electronic vote-counting systems are working correctly; to discover and recover from major failures of the electronic vote-counting software; to prevent and deter large-scale vote fraud; to provide transparency; and to give election observers evidence that the election was performed correctly. If done right, these audits provide us with a powerful defense against errors and election fraud: the paper records are a cross-check on the electronic records, and the electronic records are a cross-check on the paper. It is for these reasons that I recommend routine audits be used across the board, for both DREs and optically scanned paper ballots.

5. Dr. Wagner, why is inspection of machine software and hardware not sufficient for trusting a voting system?

As explained in my response to Question 3, it is beyond the state of the art to verify through inspection that the machine software and hardware will work correctly on election day. Given the current state of technology, it is not feasible to verify that the machine software and hardware is free of malicious logic, nor is it feasible to verify that the machine software and hardware is free of defects, flaws, and bugs.

Modern voting software and hardware is too complex to inspect completely. The software in a typical voting machine might contain hundreds of thousands of lines of source code. If all of this source code were to be printed on paper, it would fill thousands of sheets of paper. Each line of source code would have to be inspected manually by software experts, and these experts would have to understand how those lines of source code might interact with each other. This task is too complex to perform with 100% confidence; it is simply too easy to miss problems.

The U.S. Tax Code might provide a useful analogy[21]. The tax code also contains thousands of pages of material, and probably no one person understands it in its entirety. The tax code is infamous for containing loopholes that aren't obvious on first inspection; so, too, can source code contain malicious code or defects that aren't obvious on first inspection. At the same time, tax code is written to be interpreted by human judges, who might apply some degree of common sense from time to time; in comparison, software is executed by computers, who are unfailingly literal-minded, so while small ambiguities in the tax code might be minor, small ambiguities in software can be catastrophic. The analogy to the tax code is decidedly imperfect, but it might help provide some intuition about why inspection of voting software and hardware is not sufficient to trust a voting system, given the current state of technology.

A second difficulty is that, given current practice, it is difficult to be sure that the software and hardware that is running on the machine on election day is the same as what has been inspected. The existing technology does not provide any way to verify what software is running on the voting machine. Moreover, some machines have known security vulnerabilities that could allow an attacker to modify the software installed on the machine, so that the software executed on election day differs from the software that was inspected and certified. Also, there have been documented cases where uncertified versions of software were inappropriately installed and used in elections[22] [23] [24] [25].

At the same time, despite these limitations, inspection does have benefits. While it is not sufficient on its own to provide a basis for trust in voting systems, inspection—if done right—is still a good idea that can help reduce the number of voting system failures. Unfortunately, today's voting systems are not currently subject to any meaningful form of inspection by independent parties. The source code is kept secret by vendors, and access is tightly restricted. The federal testing labs—one of the few parties who are routinely given access to voting source code—do not perform meaningful inspections of source code. (The limited inspection that federal testing labs perform is more analogous to running a spell-checker on a student essay than to checking whether the writing in the essay is grammatical, coherent, meaningful, or persuasive.) In the few cases where independent experts have had the chance to inspect voting source code, they have often found serious flaws in these products which the testing labs overlooked[26]. Consequently, I believe that broader inspections of voting system software and hardware would help improve the reliability and security of elections, even though they are not on their own sufficient and would need to be supplemented with voter-verified paper records and routine manual audits.

Notes

[1] David Wagner, Written testimony before U.S. House of Representatives at joint hearing of the Committee on Science and Committee on House Administration, July 19, 2006.

[2] "Public Comment on the 2005 Voluntary Voting System Guidelines", ACCURATE Center, submitted to the United States Election Assistance Commission, September 2005.

[3] Douglas W. Jones, "Voting System Transparency and Security: The need for standard models", written testimony before the EAC Technical Guidelines Development Committee, September 20, 2004. `http://www.cs.uiowa.edu/~jones/voting/nist2004.shtml`

[4] Peter G. Neumann, Written testimony before the California Senate Elections Committee, February 8, 2006. `http://www.csl.sri.com/neumann/calsen06.pdf`

[5] Aviel D. Rubin, Written testimony before the Election Assistance Commission, May 5, 2005. `http://avirubin.com/eac.pdf`

[6] Tadayoshi Kohno, Adam Stubblefield, Aviel D. Rubin, and Dan S. Wallach, "Analysis of an Electronic Voting System", May, 2004.

[7] RABA Innovative Solution Cell, "Trusted Agent Report: Diebold AccuVote-TS System", January 20, 2004.

[8] Harri Hursti, Black Box Voting, "Critical Security Issues with Diebold Optical Scan", July 4, 2005.

[9] "Security Analysis of the Diebold AccuBasic Interpreter", Report of the California Secretary of State's Voting Systems Technology Assessment Advisory Board, February 14, 2006.

[10] Harri Hursti, Black Box Voting, "Critical Security Issues with Diebold TSx", May 11, 2006.

[11] Douglas W. Jones, "Connecting Work on Threat Analysis to the Real World", June 8, 2006.

[12] "The Machinery of Democracy: Protecting Elections in an Electronic World", Brennan Center Task Force on Voting System Security, June 27, 2006.

[13] New Yorkers for Verified Voting, "Analysis of Acquisition Costs of DRE and Precinct Based Optical Scan Voting Equipment for New York State", April 13, 2005. `http://www.nyvv.org/doc/AcquisitionCostDREvOptScanNYS.pdf`

[14] Barbara Simons, "Electronic voting systems: the good, the bad, and the stupid", ACM Queue 2(7), October 2004.

[15] Justin Moore, "How Effective is an Occasionally-Used Paper Ballot?". `http://www.cs.duke.edu/~justin/voting/paper_effectiveness.pdf`

[16] 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 2(1), January/February 2004, pp.32-37.

[17] David L. Dill, Bruce Schneier, Barbara Simons, "Viewpoint: Voting and technology: who gets to count your vote?", CACM, 46(8), August 2003.

[18] Douglas W. Jones, "Auditing Elections", Communications of the Association for Computing Machinery 47(10), October 2004, pp.46-50.

[19] Aviel D. Rubin, Written testimony before the Election Assistance Commission, June 30, 2005. `http://avirubin.com/vote/eac2.pdf`

[20] Roy G. Saltman, "Final Project Report: Effective Use of Computing Technology in Vote-Tallying", NBSIR 75-687, prepared for the Clearinghouse on Election Administration, May 1975.

[21] This analogy is taken from Barbara Simons, Jim Horning, "Risks of technology-oblivious policy", CACM 48(9), Sept. 2005.

[22] "Staff Report on the Investigation of Diebold Election Systems, Inc.", Presented before the California Voting Systems and Procedures Panel, April 20, 2004. `http://www.openvotingconsortium.org/files/shelly_diebold_report_april20_final.pdf`

[23] "Phase II County Voting System Review", R&G Associates, April 19, 2004. `http://web.archive.org/web/20041108230726/http://www.ss.ca.gov/elections/ks_dre_papers/rg_phase_II_revised_report.pdf`

[24] "E-Voting Undermined by Sloppiness", Kim Zetter, Wired News, December 17, 2003. `http://www.wired.com/news/evote/0,2645,61637,00.html`

[25] "Diebold: Voting machine maker dinged in CA: Auditor says software wasn't approved", Elise Ackerman, Mercury News, December 17, 2003.

[26] Douglas W. Jones, "Misassessment of Security in Computer-Based Election Systems", Cryptobytes 7(2), Fall 2004, pp.9-13.