# Cyber Threats to Election Integrity[1]

Herbert Lin
Stanford University
herblin@stanford.edu

## 1. THE ELECTORAL INFRASTRUCTURE

The technical infrastructure that supports elections includes systems for vote-counting and systems for voter registration. Once paper-based, both types of systems have become increasingly computerized over the last few decades.

Systems for vote counting include:

- electronic voting systems that record ballots cast by citizens in person at individual precincts.
- tabulation systems that record absentee ballots via postal mail.
- Aggregation systems that total vote counts at levels higher than precinct (e.g., systems that total all precinct totals in a given county).

Voter registration databases store information about who is eligible to vote. Proper administration of such databases entails a number of large-scale tasks that include:

- Keeping individuals who are properly registered to vote on the voter registration lists and ensuring that all of the information regarding their status is correct.
- Striking individuals from the registration lists who are not eligible to vote (e.g., individuals who have moved out of the jurisdiction).
- Ensuring that precinct-by-precinct voter registration lists are delivered to the individual precincts where in-person voting occurs (creating and delivering poll books).

[1] [The discussion of this section draws heavily on two reports from the National Research Council: *Asking the Right Questions About Electronic Voting*, Richard Celeste, Dick Thornburgh, and Herbert Lin, editors, 2004; and *Improving State Voter Registration Databases*, 2010]

Both vote counting systems and voter registration systems are part of a larger entity whose stakeholders involve various people and organizations. Political parties and candidate campaigns, the news media (both traditional and non-traditional), poll workers, pollsters, and citizens are all interested in election outcomes. Partisan stakeholders—by definition, those who want their particular candidates to win—generally seek favorable arrangements for election rules that cover everything from when and how voters vote to technology acquisition contracts. For purposes of this paper, security issues related to these other stakeholders are not addressed, though they too have important security interests.

## 2. THE NEED FOR TRUST TO BE EARNED

Measures to safeguard election integrity must be sufficiently transparent that the losers of an election are willing to accept election losses as reflecting the properly counted vote of the people in an election. Such a trust in electoral processes must be earned by those running an election, and the burden of proof is properly on election administrators to demonstrate the legitimacy of any given election result even in the face of partisan skepticism about the outcome.

As the National Research Council put it in a still-relevant report of 2005:

> **[T]rusted election processes should be regarded as the gold standard of election administration,** where a trusted election process is one that works, that can be shown to have worked after the election has been held, that can be shown to have not been manipulated and to have not led to a large number of mistaken or lost votes, and that can be shown to reflect the intent of the voters. Trusted election processes increase the likelihood that elections will be regarded as fair, even by the losing side and even in a partisan political environment.

Because so much of the technology infrastructure of elections is computer-based, applying the lessons of good cybersecurity practice (both technical and procedural lessons) is an essential element in assuring the populace than an election was conducted fairly. On the other hand, applying these lessons to elections is vastly complicated by three challenges.

First, voters have a right to cast their ballots in secret, and this right cannot be compromised by any procedure designed to audit the results of an election. Indeed, imagine difficulties of developing a financial audit procedure for a bank in which the particular monetary transactions of customers could not be associated with specific customers.

Second, an election must produce a winner even when only a few votes separate winner and loser. Small manipulations are inherently harder to detect than large ones, and at the same time are easier to perpetrate, and so the risk of election fraud is greatest when the electorate is more or less evenly divided. (And it should not escape anyone's notice that a more-or-less evenly divided electorate is exactly what characterizes much of the political landscape in the United States today.)

Third, the value of cybersecurity measures in many other computer-based systems can often be justified in cost-benefit terms, e.g., by comparing the cost of a particular security measure to the expected loss if the measure is not taken. But in the electoral context of a democracy, how does one measure the value of a vote? The reality is that budgets are finite, and election administrators have to make choices about what measures can be taken that are reasonable and adequate to address concerns about election integrity.

We now turn to the security issues of the electoral infrastructure. Because many of the security issues are common to all computerized elements of the infrastructure, they are discussed first. Security issues relevant to individual components are discussed after that.

## 3. CYBERSECURITY CONSIDERATIONS FOR THE ENTIRE ELECTORAL INFRASTRUCTURE

It is a fundamental truism about information technology that system testing can identify defects in a system (e.g., security vulnerabilities, software bugs) but no reasonable amount of system testing can prove that the system is free of defects. Testing offers evidence that a system does meet certain requirements (e.g., produces certain outputs when given certain inputs), but it is impossible to demonstrate that the system will never do something else in addition that is bad.

Deployments of a system for actual use further complicates security. System certification and testing can only evaluate the software and hardware that the vendor presents, and it does not—indeed, it cannot—take into account how the system is actually operated and maintained when it is in use. As is true for all software, bugs and therefore vulnerabilities are inevitably discovered after initial deployment, which means a system that repairs the vulnerabilities known today may well be vulnerable tomorrow. And people will make mistakes in using the technology that are not anticipated in testing.

This essential point is that assurance of security is necessarily a ongoing process that searches for vulnerabilities proactively and fixes them immediately. This point in turn undergirds two necessary elements of electoral cybersecurity.

First, security by checklist compliance (which is all that the certification process requires) does provide a baseline level of security that generally exceeds that provided by having no standards for security at all. Yet, it is known to be inferior to security assessed through an adversarial process. The best such process is an independent white-hat attack,[2] that is, a test attack by white-hat teams that do everything real attackers would do in an actual attack, taking advantage of technological or procedural flaws in the system's security posture or flaws in the human infrastructure in which the technology is embedded. The security flaws uncovered by white-hat attacks are then forwarded to responsible parties for fixing.

[2] https://www.nap.edu/catalog/10274/cybersecurity-today-and-tomorrow-pay-now-or-pay-later.

Another type of adversarial process is an independent examination of the physical hardware and software (i.e., the source code) of the system in question. Such examination will yield information about the system's ability to resist attack, and only by inspection is there any chance for discovering the potential for bad behavior when the system is actually used. Note that vendors generally resist third party inspection of source code on the grounds that allowing such access compromises their intellectual property interests in that software.

Second, "security by obscurity" is a poor security practice that is all too often employed. Security by obscurity is the practice of hiding vulnerabilities from potential attackers.[3] By contrast, disclosure of vulnerabilities provides strong incentives for system owners to fix them.

Two other shibboleths about cybersecurity also need to be addressed here. Many people believe that security of a computer is assured by not connecting the computer to the Internet. Although it is certainly true that many attacks on computer systems are delivered through the Internet, the lack of an Internet connection does not in any way guarantee security. The computer can be compromised long before it is ever delivered to the user; the user can compromise it, either wittingly or unwittingly; system updates must be applied in some way, and the updates could be compromised.

Second, cybersecurity is not just a technical problem. Human vulnerabilities can be exploited, and human beings are intimately involved in all electoral operations. Looking at the resilience of the electoral infrastructure as only or even primarily a technical issue is a profound mistake, and many of the most harmful attacks on computer systems originate with an attacker targeting a human being.

[3] More precisely, concealing the internal operation of a system does provide a layer of protection for a system. But because concealment does not actually fix vulnerabilities, these vulnerabilities can be exploited and generally such exploitation outweighs the advantages provided by obscurity.

## 4. CYBERSECURITY CONSIDERATIONS FOR VOTE-COUNTING SYSTEMS

As for the security of vote-counting systems specifically, a number of independent research efforts have demonstrated the ease with which individual electronic voting stations can be compromised with the paltry resources available to university research teams.[4] A hostile nation-state would be able to deploy orders of magnitude more resources to this task; in addition, they would have no qualms about attacking vulnerabilities that would yield higher leverage—that is, vulnerabilities that are undoubtedly present at levels higher than that of the individual machine, e.g., at the vendor level.

Addressing the security afforded by a given system is the first step assuring overall election security. A second step is ensuring that the system that has been properly certified is actually the system deployed for use by voters and not one that has been tampered with after certification. The vendor must load software onto voting machines before they are shipped to precincts. Voting machines usually sit in storage after delivery until they are pulled for deployment and use by voters. Both such stages provides tampering opportunities.

Communicating results from individual polling stations to a central tabulation authority is a third step that additional security issues—a secure process must be established to ensure that the ballot totals received at the central tabulation authority match those recorded at the precinct level. Such communication can be performed manually, by electronic transmission (e.g., over the internet or a phone line), or by physically carrying computer-readable media containing precinct-level vote totals to the physical location of the tabulation authority. Each of these methods has its security risks, but they can be mitigated with proper attention (and using more than one in parallel has security advantages as well).

[4] See, for example, https://www.ajc.com/news/state--regional-govt--politics/how-hack-elections-georgia-electronic-voting-machines/K4s5F935330BS6fGDm3CVI/.

## 5. CYBERSECURITY CONSIDERATIONS FOR VOTER REGISTRATION SYSTEMS

Voter registration systems are different from vote-counting systems in a number of ways that affect its security posture.

Perhaps the most important difference is that because voter registration systems tend to be more centralized, and thus an attacker can gain greater leverage by compromising a few voter registration systems as a way to manipulate an election rather than many voting machines. For this reason, it is plausible that voter registration systems will be subject to more sophisticated attacks than vote-counting systems.

A second important point is that for the purpose of confirming continued eligibility for voter registration, some voter registration systems draw information from other databases, such as departments of motor vehicles, departments of correction, and departments of vital statistics. Compromises in those databases (i.e., alteration or erasure of key data) could have ripple effects on the accuracy of voter registration databases (e.g., eligible voters could be purged inappropriately). Thus, security issues with these other databases (and attacks on them) could adversely affect the integrity of voter registration databases.

Third, voter registration systems entail relatively straightforward computerized functionality that is present in many commercial database systems. Thus, the underlying software of voter registration systems is likely to have been more proven through applications to multiple problem domains and exercised repeatedly; by contrast, vote-counting systems are more of a niche product that is put into operational use only sporadically.

## 6. POLICY MEASURES TO ENHANCE THE RESILIENCE OF THE ELECTORAL INFRASTRUCTURE

Given the cybersecurity issues inherent in computerized vote-counting systems, a number of measures should be taken to address them.

- All electronic vote-counting systems must have the capability for providing a voter-verified paper audit trail (VVPAT). Additionally, a percentage of electronic vote counts should be audited using the VVPAT depending on the margin that separates winning and losing candidates—this percentage would be small (but not zero) if the margin is large, and would be close to or equal to 100% if the margin were sufficiently small (e.g., 0.1% of the vote).
- The security of computerized systems in the electoral infrastructure should be addressed adversarially. This would call for a combination of white-hat attacks and independent code inspection. Concerns about intellectual property protection are understandable, but at root is the point that in the absence of independent code inspection, the result is an important instrument of public policy that is not public or at least subject to independent scrutiny. Legitimate concerns can be addressed through the use of carefully crafted non-disclosure and/or non-compete agreements. The former would permit public discussion of security flaws found but also specify details that could not be disclosed. Non-compete agreements could be used to provide vendors with assurances that allowing code inspection would not enable those viewing code to become competitors. Reports derived from white-hat attacks and code inspection would report their findings publicly, so as to put pressure on vendors and election administrators to fix the problems. Bug bounties should also be offered in searching for vulnerabilities.

Certain other changes in the organizations surrounding election infrastructure would also enhance its resilience and increase public confidence in the conduct of an election.

- Election administrators should implement training programs for themselves and their staffs to implement basic cybersecurity practices.[5]
- Elections should be administered by nonpartisan officials.

[5] For example, the Kennedy School of Government has published a Campaign Cybersecurity Playbook that describes what political campaigns should do regarding cybersecurity. Much of that playbook is applicable to those who work on election administration. See https://www.belfercenter.org/sites/default/files/files/publication/CampaignPlaybook_0.pdf.

- Vote-counting and voter registration systems acquired from vendors whose senior leadership are demonstrably partisan should be scrutinized with extra care and attention.[6]
- Voters should be able to cast votes in person over an extended period of time (several days). Apart from all of the reasons that early voting make sense from the standpoint of voter convenience, it is simply a reality that technical problems often appear in complex computer-based equipment when placed into widespread operation with real users, and attempting to deploy fixes on a time scale of hours is often not feasible.

Finally, efforts to enhance the resilience of electoral infrastructure—both technical and organizational—will have to be continual. Accordingly, a regular funding stream must be made available to electoral administrators for cybersecurity purposes—by the nature of the threat and of technology, cybersecurity is not a problem that can be solved once and for all.

[6] As one illustration of such a partisan leaning, the CEO of a vendor trying to sell voting machines in Ohio in 2003 said that he was "committed to helping Ohio deliver its electoral votes to the president next year." See Julie Carr Smyth, "Voting Machine Controversy," *Cleveland Plain Dealer*, August 28, 2003. Reprinted at https://www.commondreams.org/headlines03/0828-08.htm.