Quantitative Security Analysis of Internet Voting vs. Two Other Voting Systems

A position paper for the "Workshop on UOCAVA Remote Voting Systems"
Aug 6-7, 2010.

Eric L. Lazarus, David L. Dill, Bruce Schneier

We report on preliminary results of a quantitative security analysis comparing three different voting systems: Hand-Counted Paper Ballots (HCPB), Remote Poll-site Voting with Ballot-on-Demand (RPVBoD), and Internet voting. We chose HCPB to establish a security baseline, since it is a simple, familiar, and relatively well-understood system. RPVBoD is a new proposal for overseas military voting where jurisdiction-specific paper ballots are printed just before the voter votes, the ballots are marked by the voters, and the paper ballots are returned the voters' local jurisdictions to be counted. By Internet voting, we mean a system where voted ballots are returned electronically via the public Internet.

We chose the cost of an attack to be *minimum attack team size* – the number of *knowing participants required to conduct a successful attack that reverses the outcome of the election*. A system that can be compromised by a single attacker is much less secure than a system that requires hundreds of cooperating attackers. Although arguments can be made for other cost metrics, attack team size has several advantages: the risk of detection is generally proportional to the number of participants in an attack; other costs, including monetary, will be aligned with the difficulty of recruiting people who will knowingly commit election fraud and placing them in positions where they have the necessary access to attack the system; attackers could be extremely well-funded, so money is not as scarce a resource as attack team members; and it is easier to estimate attack team size than other metrics we have considered.

The analysis revealed that *stealing a well-run election conducted with HCPB in a U.S. Senate election in a medium-size state requires roughly 1,272 people to engage knowingly in miscounting votes, while stealing an Internet election can be performed by a very small team, perhaps even a single individual. A successful attack against an RPVBoD election requires an team size of at least 114, even in a close election.* These numbers are necessarily approximate, but they reliably show that **Internet voting is much less secure than other voting systems, including a realistic alternative system for overseas voting.**

More generally, we believe that systematic threat analysis with a simple quantitative metric could greatly aid policy decisions about voting systems because it provides a structure for organizing the many complex issues involved in assessing system security. In particular, when results are controversial, a good analysis exposes the specific assumptions that are the sources of disagreement, significantly improving the quality of discussion and allowing for targeted security improvements.

Below, we briefly describe the analysis method, and then describe some of the specific attacks for each voting system that illustrate some of the security issues.

Method

Attack team size calculations are computed based on a variety of factors including the number of precincts, the number of voters (both total registered and votes cast), and margin of victory. To arrive at specific numbers, we analyzed the United States Senate election in Virginia in 2006, and assumed conservatively that an attacker would try to change 1%, or 23,000, of the votes to minimize the risk of detection while having a chance of changing the outcome of a close election (the margin of victory in this particular race was about 0.5%). Virginia is estimated to have 124,689 overseas eligible voters.[1]

We assumed that the administration of each voting system follows the best feasible practices. If a voting system is inadequate because of suboptimal practices, the public interest is better served by improving the practices, not by choosing a different voting system. Also, assuming less than best practices fails to distinguish between voting systems that have the potential to be secure and those that are hopeless even with best practices.

The analysis of the HCPB system assumed that voters vote in polling places on paper ballots, and that the ballots are counted by hand in public view with observers present.

RPVBoD is a new proposal, so we describe the system and assumptions in a bit more detail. In this system, voters travel to polling places where the appropriate ballot for their jurisdiction is printed when they sign in, except for voters who are not realistically able to visit a polling place, who mail their ballot to their home jurisdiction (we assume that these voters need to supply some reason for voting outside of the polling place, and that they are not numerous). Voters mark their ballots by hand. Ballots are placed in envelopes by the voter at the polling place, and are deposited by the voter into a secure ballot box in public view. The envelopes are provided at the polling place with mailing labels that are printed on demand at the same time as the ballots, with addresses and postal bar codes provided by the voter's home election office. The ballot boxes are delivered by military transport, with guards chosen at random from different units and political parties, to centers where the ballots can be transferred to the civilian postal service, and where that transfer can conveniently be observed by members of the public and political parties. The paper ballots are then sorted and shipped to the individual jurisdictions for counting with other ballots. As a defense against wholesale ballot stuffing, we assumed that, for each polling place, the number of voters for each county are published immediately after the election.

For the analysis of the Internet voting system, we assumed that voted ballots are returned electronically over the public Internet by voters. We assumed the use of reasonable best practices for computer and Internet security.

1 This estimate is from Dr. Michael McDonald at George Mason University, based on statistics from the Federal Voting Assistance Program and other sources (http://elections.gmu.edu/Turnout_2008G.html).

Key attacks against Hand-Counted Paper Ballots (HCPB)

The analysis in this case was to provide a baseline for comparison with the other methods. It is appropriate for voting in domestic polling places, not remote voting. The HCPB election system has been refined over centuries, so the issues are well understood (even if best practices in administering such elections are not always followed). We assume that there are controls in place to guard against confusing ballot design and misallocation of polling place resources, that poll workers and ballot counters are assigned randomly, and that voting and counting of votes are publicly observed, that ballots are hand-counted in each precinct immediately after the close of the polls and the results published immediately after completion of the count, that the lists of voters who signed in are compared with the numbers of votes cast, etc. Given assumptions, the most effective attack has a relatively large attack team size:

Subvert the hand count: Hand-count teams could deliberately count incorrectly. This will typically require the poll workers plus the supplemental counting teams. In the example Senate race, that would be require all the poll workers and counters at 53 polling places. We calculate the number of polling places that would need to be attacked by computing the number of people voting at each precinct on average, based on the number of precincts in Virginia, times a maximum percentage of stolen ballots, 15%, intended to minimize the risk of detection. We multiply by 4 pollworkers and 20 hand counters per polling place for a total attack team size of 1,272. (In theory, this number could be raised substantially if the assignment of counters and poll workers were done in a transparently random way so that the attacker would be forced to subvert nearly the entire pool of counters.)

Key attacks against Remote Poll-site Voting with Ballot-on-Demand (RPVBoD)

Ballot stuffing in transit: This class of attacks involves adding to the collection of ballots as they are being transported from the polling sites to the voters' jurisdictions. For example, ballot transport guards open ballot boxes and add counterfeit ballots and re-seal them.

We based the analysis on an attacker goal to add 23,000 votes. We assumed guard teams of six people, assigned at random, and that each guard team transports no more than 10,000 ballots and that they would not steal more than 25% of the ballots to minimize the risk of detection. That would require 10 teams, for a total of 60 guards. Also, since we have assumed that each polling site publishes the number of voters from each county, that data would have to be altered. We assume that there are no more than 5,000 voters per polling site. If there are 6 poll workers per polling site, this would require 120 poll workers to participate in the attack, so the total attack team size would be 180.

Ballot box stuffing by voting center staff members: This is similar to the previous class of attacks, but involves adding ballots before they leave the polling site. Poll workers could identify legitimately registered voters who choose not to vote and create votes signed by these

voters. This attack requires participation by the staff at attacked polling site. It is made more difficult by the random selection of poll workers.

The targeted number of votes to steal by this method is 23,000. As before, we assume no more than 5,000 votes per polling site and that no more than 25% of the ballots could be stolen without arousing too much suspicion, so each team can steal appropriately 1,250 votes. These assumptions lead to a requirement that at least 19 teams, totaling 114 staff be involved.

In-polling-place voter impersonation: This attack requires individuals to pretend to be legitimate voters without the support of corrupt poll workers. In the military context this involves creating fraudulent IDs and identifying people willing to vote with them. In the context of highly controlled overseas installations, this will likely require at least one "subversion coordinator" per polling place to coordinate fake ID distribution and fraudulent voter recruitment, and the defeat of ID check processes. Of course, there will be some polling place contexts which "everybody knows everybody" in which this sort of attack will not work at all.

As before, the goal is to steal 23,000 votes. If we assume that one person can only impersonate 10 voters without unacceptable risk of detection, the result is an attack team size of about 2,300, leading to the conclusion that this is not a particularly effective means of attack.

Key attacks against Internet voting

Unfortunately, there are many attacks on Internet voting with a small attack team size. Given limited space, we chose not to provide a comprehensive list, but to focus on a few of the more interesting attacks. For example, we do not discuss in detail attacks that direct voters to the wrong vote servers, either by hijacking parts of the Internet infrastructure or by deceiving voters ("phishing" attacks). We also omit attacks based on the surprising difficulty of authenticating the identity of voters. Both of these classes of attacks have very small attack team sizes.

Malware on voting terminal: Malicious software (malware) such as viruses on the PC used by the voter would have complete control over the voting process (the malware can observe everything use user types and sees). Hence, malware can alter votes, destroy votes, steal passwords and pins, make it impossible to vote, violate voter privacy, or trick the voter in many ways. This class of attack is shown to be feasible every day by the existence of so-called "botnets" consisting of millions of virus-infected PCs controlled by cybercriminals, unbeknownst to their owners.[2] This attack requires only a single individual to design, develop and deploy the malware.

Server Subversion: Vote data on the vote servers can be modified manually or by malware. Malware could be inserted by members of the original development team, by staff within the local elections office or by external hackers. Malware could also be inserted into the embedded

2 The"Mariposa" botnet, which was taken down by Spanish police in early 2010, was estimated to have had almost 13 million PCs ("Spanish police take down massive Mariposa botnet," IDG News Service, March 2, 2010).

firmware of hardware components of the server, such as disk drives or motherboards, by attackers involved in their design and/or manufacture. Election office staff or external hackers could edit vote data. The skills to conduct these attacks are high in some cases, but also widely available in cybercrime syndicates and in national security organizations, and the attack team size could be as low as a single person.

Coercion/Vote buying: Internet voting from personal computers creates new opportunities for computer-assisted vote-buying or coercion. The enabling condition for coercion and vote-buying is that the voter be able to prove how he or she voted to someone else, who can pay for a desired vote or threaten to punish the voter for an undesired vote. Voters could be asked to install software that monitors their vote, or vote through a "man-in-the-middle" proxy that can inspect the vote before submitting it to the actual vote server. Incentives could be payment to the voter or retribution of various forms. The effectiveness of "scareware" in inducing users to pay for bogus anti-virus products shows that even empty threats may cause some voters to change their votes. Although absentee voting is also generally vulnerable to vote buying and coercion on a small scale, Internet voting creates the ability for a very small team of attackers to "look over the shoulders" of a much larger set of voters. Generally, bribing or threatening many thousands of people over the Internet would require only a very small number of instigators.

Interestingly, this attack has a large attack team size. The voters being bribed or threatened are "knowing participants" in the attack, so the number of attackers is formally about 23,000, even though there may be just one instigator. The large attack team size reflects a realistic probability that several voters would come forward with information about the attempts to coerce them. Nevertheless, we believe this attack needs to be considered because it is unclear what legal remedies would be available even if the attack were proven to have occurred.