

# **A Security Analysis of the Secure Electronic Registration and Voting Experiment (SERVE)**

January 5, 2004

Dr. David Jefferson, [d\\_jefferson@yahoo.com](mailto:d_jefferson@yahoo.com)

Dr. Aviel D. Rubin, [rubin@jhu.edu](mailto:rubin@jhu.edu)

Dr. Barbara Simons, [simons@acm.org](mailto:simons@acm.org)

Dr. David Wagner, [daw@cs.berkeley.edu](mailto:daw@cs.berkeley.edu)

## Executive Summary

This report is a review and critique of computer and communication security issues in the SERVE voting system (Secure Electronic Registration and Voting Experiment), an Internet-based voting system being built for the U.S. Department of Defense's FVAP (Federal Voting Assistance Program). The authors are members of SPRG (the Security Peer Review Group), a panel of experts in computerized election security that was assembled by FVAP to help evaluate SERVE. Our task was to identify potential vulnerabilities the system might have to various kinds of cyber-attack, to evaluate the degrees of risk they represent to the integrity of an election, and to make recommendations about how to mitigate or eliminate those risks.

The SERVE system is planned for deployment in the 2004 primary and general elections, and will allow the eligible voters first to register to vote in their home districts, and then to vote, entirely electronically via the Internet, from anywhere in the world. Besides being restricted to overseas voters and military personnel, SERVE is currently limited to people who vote in one of 51 counties in the seven states (Arkansas, Florida, Hawaii, North Carolina, South Carolina, Utah, and Washington) that are participating. The program is expected to handle up to 100,000 votes over the course of the year, including both the primaries and the general election. (By comparison, approximately 100 million votes were cast in the 2000 general election.) The eventual goal of SERVE is to support the entire population of eligible overseas citizens plus military and dependents. This population is estimated to number about 6 million, so the 2004 SERVE deployment must be judged as a prototype for a very large possible future system.

Our conclusions are summarized as follows:

- a) The SERVE system has numerous fundamental security problems that leave it vulnerable to a variety of well-known cyber attacks, any one of which could be catastrophic.
- b) Such attacks could occur on a large scale, and could be launched by anyone from a disaffected lone individual to a well-financed enemy agency outside the reach of U.S. law. These attacks could result in large-scale, selective voter disenfranchisement, or privacy violation, or vote buying and selling, or vote switching even to the extent of reversing the outcome of many elections at once, including the presidential election. With care in the design of the attacks, some of them could succeed and yet go completely undetected. Even if detected and neutralized, such attacks could have a devastating effect on public confidence in elections.
- c) It is impossible to estimate the probability of a successful cyber-attack (or multiple successful attacks) on any one election. But we show that the attacks we are most concerned about are quite easy to perpetrate. In some cases there are kits readily available on the Internet that could be modified or used directly for attacking an election. And we must consider the obvious fact that a U.S. general election offers one of the most tempting targets for cyber-attack in the history of the Internet, whether attacker's motive is overtly political or simply self-aggrandizement.
- d) The vulnerabilities we describe cannot be fixed by design changes or bug fixes to SERVE.. These vulnerabilities are fundamental in the architecture of the Internet and of the PC hardware and software that is ubiquitous today. They cannot be eliminated for the foreseeable future; indeed, if they were to be eliminated, it would be only as a result of wholesale redesign and replacement of much of the hardware and software security systems that are part of, or connected to, today's Internet.
- e) We have examined numerous variations on SERVE in an attempt to recommend an alternative system that might deliver somewhat less voter convenience in exchange for fewer or milder security vulnerabilities. However, all such variations suffer from the same kinds of fundamental vulnerabilities that SERVE does; regrettably, we cannot recommend any of them.
- f) The SERVE system might appear to work flawlessly in 2004, with no successful attacks detected. Unfortunately, a seemingly successful voting experiment in a U.S. presidential election involving seven states would be viewed by most people as strong evidence that SERVE is a reliable, robust, and secure voting system. Such an outcome would encourage expansion of the program by FVAP

in future elections, or the marketing of the same voting system by vendors to jurisdictions all over the United States, and other countries as well.

However, the fact that no successful attack is detected does not mean that none occurred. Many attacks, especially if cleverly hidden, would be extremely difficult to detect, even in cases when they change the outcome of a major election. Furthermore, the lack of a successful attack in 2004 does not mean that successful attacks would be less likely to happen in the future; quite the contrary, future attacks would be more likely, both because there is more time to prepare the attack, and because expanded use of SERVE or similar systems would make the prize more valuable. In other words, a “successful” trial of SERVE in 2004 is the top of a slippery slope toward even more vulnerable systems in the future.

- g) In spite of the fact that we recognize the need to offer better support for voting for overseas Americans, we are forced to conclude that the best course to take is not to field the SERVE system at all. Because the danger of successful, large-scale attacks is so great, we reluctantly recommend shutting down the development of SERVE and not attempting anything like it in the future until both the Internet and the world’s home computer infrastructure have been fundamentally redesigned. In Appendix C of this report, we offer an alternative to SERVE that we believe is much more secure.

## 1. Introduction

This report is a review and critique of computer and communication security issues in the SERVE voting system (Secure Electronic Registration and Voting Experiment), an Internet-based voting system being built by Accenture and its subcontractors for the U.S. Department of Defense's FVAP (Federal Voting Assistance Program). The authors are members of SPRG (the Security Peer Review Group), a panel of experts in computerized election security that was assembled by FVAP to help evaluate SERVE. Our task was to identify potential vulnerabilities the system might have to various kinds of cyber-attack, to evaluate the degrees of risk they represent to the integrity of an election, and to make recommendations about how to mitigate or eliminate those risks.

The analysis and conclusions outlined here are based on two three-day meetings we had with the FVAP sponsors and the primary technical architects of SERVE; these were held in July, 2003 at Caltech in Pasadena, California, and in November, 2003 at Accenture in Reston, Virginia. Many issues and design improvements were proposed and accepted in those meetings; in this report we concentrate on the remaining security issues that were not resolved.

### 1.1 What is SERVE?

Part of the mission of FVAP is to reduce the barriers to registration and voting for two groups of eligible voters: (1) American citizens living outside the U.S., and (2) military personnel and their dependents, regardless of where they reside in the U.S. or overseas. For Americans living overseas, voting can be a daunting task; it can take five or more trips through U.S. and foreign mail services to request voter registration forms and absentee ballots from the home county, to receive them, and then to send them back—a process that is time-consuming and unreliable at best, and must be accomplished in timely manner to avoid missing legal deadlines. The process is so clumsy for soldiers who are mobile, or who are located where mail service is poor, that their voting rates are believed to be quite low.

The SERVE system is planned for deployment in the 2004 primary and general elections, and will allow the eligible voters both to register to vote in their home districts, and then to vote, entirely electronically via the Internet, from anywhere in the world. To participate, an eligible voter first must enroll in the SERVE program, which can be done completely electronically if the voter has suitable military ID (specifically, a Common Access Card), or by presenting suitable citizenship and ID documents face-to-face to a trusted agent, e.g. a military officer or other designated official who plays a role similar to that of a notary public. After enrollment, the voter will be able to register to vote, and then to vote, in one or two short sessions from any Internet-connected PC. The PC must run a Microsoft Windows operating system and either the Internet Explorer or Netscape web browser. The browser must be configured to enable JavaScript and either Java or ActiveX scripting, and it must also permit session cookies; however, no additional hardware or software is required.

SERVE is designed as a Web-based service. Voters connect to a central server using a standard browser, as just described. Both registration and voting are accomplished through the web interface. SERVE requires direct interaction between the voting service and the Local Election Official (LEO) in the voters' home precincts. Thus, when people register to vote, their information is sent by the central web server to the LEO, and the LEO updates its database. When someone votes in the election, the completed ballot is stored on the central server, and later downloaded by the LEO, who updates the local tally. The communication between the user's web browser and the voting application on the central server is protected using the encryption and authentication built into the Secure Socket Layer (SSL) protocol. Once that connection is established, an ActiveX control (described later) is downloaded to the client, because the voting application requires functionality that is not available in current standard browsers. For Netscape users, a Java applet runs that interprets the ActiveX. For Internet Explorer users, the ActiveX control runs natively on the machine.

Besides being restricted to overseas voters and military personnel, in 2004 SERVE will be limited to people who happen to vote in one of 51 counties in the seven states (Arkansas, Florida, Hawaii, North Carolina, South Carolina, Utah, and Washington) that have agreed to participate. The 2004 trial is expected to handle up to 100,000 votes over the course of the year, including both the primaries and the general election. (By comparison, approximately 100 million votes were cast in the 2000 general election.)

One goal is to evaluate whether SERVE might be suitable for expansion in the future to all overseas voters. The total number of eligible overseas citizens plus military and dependents is substantial, estimated at about 6 million eligible voters. Furthermore, systems similar to SERVE might eventually be offered by Accenture or other vendors for certification in many more states, and with all voters eligible to use it instead of just a limited population. Therefore, any security analysis must assume that the use of SERVE could be significantly expanded in future years.

## **1.2 Brief History of Internet Voting**

The SERVE system is a follow-on to an earlier FVAP voting system called VOI (Voting over the Internet). VOI was built by a different general contractor (Booz-Allen & Hamilton) and used a different architecture and codebase, so VOI and SERVE can be compared only in a general way. VOI was used only in the 2000 general election, handling a total of 84 votes in four states (Florida, South Carolina, Texas, and Utah). All were real votes, not test ballots.

The FVAP office issued a report on the VOI system in June, 2001 (*Voting Over the Internet Pilot Project Assessment Report*). Only a small part of the report is devoted to security issues, but most of the concerns we mention below were discussed. One conclusion of the report is that the VOI experiment was so small that it was not a likely target of any attacks, and that even a successful attack would almost certainly have made no difference in the outcome of any election. (The fact that 50 votes were cast in Florida using VOI, and that a change of 269 votes in the official tally of that state would have resulted in Al Gore becoming President, shows how dangerous such an assumption can be. We note that Florida is participating in 2004 in the SERVE program.)

The VOI report also defined away key parts of the Internet voting security problem by taking the position that “In the VOI Pilot System the citizen’s workstation was outside the security perimeter of the system”. Thus VOI made no attempt to defend against some of the most serious attacks to which it was vulnerable.

However, the VOI report did express concern about security problems with “remote” Internet voting (i.e., voting from any Internet-connected computer, anywhere in the world, as permitted under SERVE), explicitly declining to recommend remote voting until such time in the future when the most serious threats have been resolved:

“[Remote Internet voting] is subject to the same security concerns as the current VOI System. For this reason, we cannot recommend [it] as an immediate follow-on development to the VOI Pilot. .... Therefore, we recommend that research continue on these security issues so that this alternative could be implemented in the future when adequate security measures are available to counteract the malicious software (e.g., virus and Trojan Horse) threat and denial of service attempts.” (Section 6.2.4)

In spite of this recommendation, the SERVE program is deploying remote Internet voting as the follow-on to VOI, even though the malicious software and denial of service threats have not been resolved.

In 2000 there were several other experiments with Internet voting in U.S. public elections. In some cases the votes counted officially; in others they did not. The largest and most well-known was the Arizona Democratic presidential primary, conducted by election.com (now part of Accenture) in March of that year, in which approximately 85,000 votes were cast and counted. The Reform Party national primary was also conducted over the Internet that summer, as were various nonbinding Internet voting experiments in some counties of Washington, California, Arizona and elsewhere.

Several studies of Internet voting, including the security issues, were conducted in 1999-2000. The first was by the California Secretary of State’s Task Force on Internet Voting, whose report was issued in January, 2000 and is online at <http://www.ss.ca.gov/executive/ivote>. That report was the first to clearly articulate most of the technical security issues regarding Internet voting in general, and it conspicuously failed to recommend that the state push ahead with “remote” Internet voting (e.g., from PCs at home, office, schools, libraries, and cybercafes), because of the numerous security problems it detailed. Those security problems had no good solutions at the time, and now, four years later, they remain unsolved.

Another study was conducted by the Internet Policy Institute with funding from the National Science

Foundation. Their report (*Report of the National Workshop on Internet Voting: Issues and Research Agenda*, referred to as the NWI Report) was based on a conference held in October, 2000, and was published in March, 2001. The following is a key paragraph from the Executive Summary of that report:

*“Remote Internet voting systems pose significant risk to the integrity of the voting process, and should not be fielded for use in public elections until substantial technical and social science issues are addressed [italics in the original]. The security risks associated with these systems are both numerous and pervasive, and, in many cases, cannot be resolved using even today’s most sophisticated technology. In addition, many of the social science concerns regarding the effects of remote voting on the electoral process would need to be addressed before any such system could be responsibly deployed. For this reason, it is imperative that public officials educate themselves about the dangers posed by remote Internet voting, and the ramifications of failure on the legitimacy of the electoral process.”*

The “numerous and pervasive” security risks referred to in the NWI Report are similar to those described in the California task force report.

The Caltech/MIT Voting Technology Project published its report (*Voting: What Is, What Could Be*) in July, 2001. It is available online at <http://www.vote.caltech.edu/Reports/2001report.html>. That report was similarly pessimistic about the security of Internet voting, stating that “Remote Internet voting poses serious security risks. It is much too easy for one individual to disrupt an entire election and commit large-scale fraud.”

As we see, all three major studies of Internet voting in 1999-2001 concluded that Internet voting is fraught with major risks that have no short term solution. We know of no major studies anywhere, then or since, that have come to any different conclusion. Those risks remain as relevant and dangerous today as ever.

In sections 2, 3 and 4 of this paper we argue that all of the security vulnerabilities that were articulated in these 1999-2001 studies are present in the SERVE architecture, along with at least one additional risk that was not emphasized then, namely potential insider fraud. We claim that, given current technology, these vulnerabilities are inherent in any Internet voting architecture that allows people to vote from private computers, and that no technological innovation will change that in the foreseeable future.

### **1.3 Criteria for assessing the security of SERVE: How much security is enough?**

In evaluating the security of SERVE, we need a standard against which to compare it, i.e. an answer to the question “How much security is enough?” No system is perfect, but we must not allow unacceptable risks of election fraud to taint our national elections.

On the one hand, election security has to be viewed as a component of national security, since the very legitimacy of democratic government depends on elections that are fair, open, trustworthy, and seen to be so. This would argue for the very highest security standards—ideally that not a single vote be lost, forged, spoiled, miscounted, bought, sold, or otherwise compromised under any credible threat scenario, even if the attacker has lots of resources, full knowledge of the SERVE architecture and an inside confederate.

On the other hand, the SERVE system is designed to be a form of, or replacement for, absentee voting. Absentee voters vote from somewhere other than the precinct polling location, traditionally by marking or punching a paper ballot and mailing it back to the county officials, although faxing is sometimes permitted. In some western states 30% or more of all votes are absentee; Oregon in particular eliminated its precinct polling places, so all votes there are absentee. Since absentee voting is the model for SERVE, arguably that should be the security baseline against which SERVE should be compared. Our analysis is therefore premised on the following principle: *At the very least, any new form of absentee voting should be as secure as the current absentee voting system and should not introduce any new or expanded vulnerabilities into the election beyond those already present.*

Absentee voting is not as secure as precinct based voting. While absentee voting procedures offer a fair degree of security and privacy, there are some inherent vulnerabilities of which everyone is aware and that

we as a society have agreed to tolerate. There are many ways for an attacker to compromise a *small* number of absentee votes without detection, e.g. by spying on voters, or coercing them, especially in institutional situations like nursing homes; or by paying voters; or by interfering with the transport of ballots through the mail, etc. However, the key point that makes absentee voting tolerable in spite of its many vulnerabilities is that it is very difficult to mount any kind of large-scale or automated attack on the process without getting caught. There are no single points of vulnerability through which many absentee votes pass except at the offices of county election officials and their local post office where, in both cases, there are numerous procedural and legal controls so that the perpetrators are at great risk of getting caught. Even if a vote fraud scheme were perpetrated in one of those places and somehow escaped detection, its scope would be limited to a single county, and in most states to only the relatively small percentage of absentee ballots cast in that county, so that the risk and penalties for getting caught outweigh the value of the small number of votes that might be compromised.

This security level—some vulnerability to undetected small-scale attacks, but little or none to large-scale attacks—seems a reasonable goal for a new voting system such as SERVE, one that would assure that its addition to the mix of other voting systems already in use would not reduce the overall security that elections currently possess, nor add new vulnerabilities of a different character or scale. What we must avoid at all costs is any system in which it is possible for a successful large-scale or automated (computerized) attacks to compromise many votes. It must be essentially *impossible* that any such attacks might go undetected; or that such an attack might be so easy and inexpensive that a single person acting alone could carry it out; or that the perpetrator(s) of such an attack might never be identified; or that the attack might be carried out remotely, from foreign soil, possibly by an foreign agency outside the reach of U.S. law, so that the attackers face little or no risk. Any voting system with any of these vulnerabilities, we believe, is completely unsuitable for use in U.S. public elections. Unfortunately, the SERVE system has all of these vulnerabilities, as we document in the rest of this paper.

## **1.4 Voting System Threats**

Any truly democratic voting system must have ways of dealing with five important threats. A first serious threat is disenfranchisement, either of individuals or of classes of voters. A major concern is that classes of voters could be disenfranchised, based on the likelihood of their voting a particular way. In contrast to the United States' history of gradually enfranchising classes that had previously been disenfranchised, Internet voting provides opportunities for selective disenfranchisement that can be difficult to detect, and even harder to ameliorate.

A second threat is that a voter's ballot could be modified by a third party without the voter's consent. With conventional paper ballots, this could be done by, say, adding a vote for an office for which the voter had not voted or by adding extra votes so that the voter's ballot is invalidated. As we shall see, however, electronic voting creates new opportunities for ballot modification that had not previously existed.

A third threat is the loss of privacy, or the undermining of the secret ballot. It is harder to protect the voter's privacy when voting is done using either a paper or an electronic absentee ballot than when voting occurs at a polling station.

A fourth threat is that a voter may vote more than once. An individual might, for example, cast the vote of an aged or ill relative. Or a poll worker might cast ballots at the end of the day on behalf of voters who never appeared at the polls. As in the case of ballot modification, Internet voting provides new opportunities for multiple voting.

A fifth threat that is intrinsic to absentee voting is that of voter coercion, vote buying, and vote selling. As we show, this threat also is amplified by Internet voting.

Finally, a theme that all these threats have in common is the issue of scale. Computers are extremely good at automating repetitive tasks, but this cuts both ways: It means that it is also easy to automate attacks. When computer security systems fail, typically they fail on a large scale. A major risk in any centralized Internet voting scheme like SERVE is that a single failure could affect hundreds of thousands of voters, whereas current non-Internet systems are carefully designed to avoid any single point of failure that might affect large numbers of votes.

## **1.5 Problems Exacerbated by SERVE**

**Privacy.** SERVE aims to provide at least the same levels of privacy and security as conventional absentee voting. Since different states process absentee mail ballots differently, we shall discuss how California attempts to protect the privacy of the mail absentee ballot. The voter inserts the ballot in an inner envelope, which in turn is inserted into an outer envelope. Information identifying the voter, such as the voter's name and signature, is written on the outer envelope. When a mail ballot is received, the name on the outer envelope is checked against the voter list. Assuming it matches the name of a registered voter who has not yet voted, the outer envelope is opened. The timing of the opening of the inner envelope is determined by state and county regulations. While privacy concerns would dictate that the inner envelope be opened such that it cannot be linked to the outer envelope, this may not always happen. And even if the best precautions are taken, the voter who forgets to put the ballot in the inner envelope will have his/her ballot made visible as soon as the outer envelope is opened.

SERVE attempts to separate the name of the voter from his/her vote through the use of public key cryptography. In public key systems, each participant has a pair of keys consisting of a private key and a public key. The private part is known only to the participant, and, as the name implies, the public portion is available to everyone. In each voting district that is participating in SERVE, an election official (called the Local Election Official, or LEO, by SERVE) generates such a key pair. The LEO's public key is used to encrypt (scramble) the ballots of SERVE voters from that district. Once a ballot has been encrypted, it can be read only if it is decrypted through the use of the unique private key known only to the LEO.

When the voter uses SERVE to cast a ballot, his/her web browser sends the completed ballot, along with identifying information such as the voter's name, to a SERVE web server. This information is transmitted to SERVE in encrypted form, sent in a way that only the SERVE web server can decrypt it. Because the ballot is encrypted during transmission, if someone were to intercept the encrypted ballot en route, it would be impossible to decipher the actual vote, although the voter's name could be read. Note that, at this point, the LEO's key pair has not yet been used.

When the ballot is received, SERVE verifies that the voter is registered and has not yet voted. It then separates the ballot from the voter's name, encrypts the ballot using the LEO's public key (so that only the appropriate LEO will be able to decrypt this encrypted ballot), and stores the encrypted ballot for later transmission to the LEO. SERVE retains the encrypted ballot (even after a copy has been sent to the LEO). SERVE also places the voter's name on a list of people who have already cast a vote, so that they will not be allowed to vote a second time.

This architecture introduces several privacy risks. First, the LEO can request from SERVE a list of names of voters from the LEO's district who have already voted via the Internet and the list (re-ordered randomly) of encrypted ballots for those voters. If a curious LEO makes the request sufficiently often, it should be possible for the LEO to deduce how each individual voter voted, an obvious privacy risk.

Second, the fact that ballots exist unencrypted on the server for a brief period before being encrypted with the LEO's key introduces other privacy risks. For instance, the SERVE system administrators could view how people have voted. Also, if the SERVE machines were compromised, the unencrypted ballots could be revealed to unauthorized third parties. See Appendix A for further discussion of this risk.

Third, the fact that the encrypted ballots together with the voters' names are stored in a SERVE database means that anyone with access to a LEO's private key could determine the votes of all SERVE voters from that voting district, another significant privacy risk.

Fourth, the encrypted ballots are stored for a long period of time on SERVE's computers, which increases the window of risk. We understand that, at a minimum, encrypted ballots will be retained until 18 months after end of elections. We believe it is possible that encrypted ballots may remain accessible for an even longer period of time, and perhaps indefinitely, for instance on backup tapes or other computers, independent of the intention of the developers. Indeed, information often remains retrievable longer than developers intend; in modern systems, information is typically copied to so many locations that it can be challenging to find and erase them all. Consequently, it is conceivable that voter privacy could be compromised at a future date if the information ever landed in the wrong hands and old system keys were exposed. Mathematical breakthroughs in the future, for example, could expose old private keys.



**Vote Buying/Selling.** Vote selling is a problem in all elections, but it is a special concern for Internet voting, since the Internet can facilitate large scale vote buying and selling. During the 2000 presidential election we saw the first Internet based attempt at vote swapping in a presidential election with the creation of a website to facilitate vote swapping between Gore and Nader voters. While the Gore/Nader swapping depended on the honor system and no money changed hands, a similar approach could be used to provide enforced vote swapping or vote bartering services, or to purchase votes from SERVE voters. Due to the ease of automating such attacks, the deployment of SERVE could potentially lead to vote buying or swapping on a larger scale than has been seen before.

The most straightforward vote-buying scheme would involve the selling of voting credentials, namely personally identifiable information and the voter's password or private key. One possible defense we considered would be for SERVE to prohibit the submission of multiple votes from the same Internet address. That would prevent someone who collects many credentials from submitting multiple votes. Restricting the number of submissions from any particular web address is not a strong defense, however, because it is possible for a purchaser of votes to fool SERVE into thinking that the votes were coming from different addresses. Also, due to proxy servers, legitimate users often appear to come from the same IP address, so this hypothetical defense may not be deployable in practice. An extreme example of this is AOL, which uses the same IP address for all users coming from its domain.

Another approach to vote-buying would be for the buyer to provide the seller with a modified version of the ActiveX component being used by SERVE. An appropriately modified version could ensure that the voting is done according to the wishes of the vote purchaser. There does not seem to be any way for SERVE to defend against this style of vote buying.

A vote buyer could also rely on an honor system similar to the one used in the Gore/Nader swap on the assumption that the majority of those selling their votes will be "honest".

**Coercion.** Coercion is a problem with all forms of absentee voting, since the voter is not guaranteed the privacy of the voting booth. But it is even more of a problem with Internet voting, especially if the voter is not using his/her machine, since the owner of the machine may have installed software that would record what the voter is doing. While military regulations prohibit coercion by superior officers, the risk remains.

**Large Scale Impact.** Because SERVE is vulnerable to many different types of attacks, a significant percentage of votes cast over the Internet are also vulnerable, and a single successful attack might be able to affect a large fraction of all votes cast through SERVE. By contrast, when voting is conducted at physical precincts on mechanical devices or with paper ballots, vote manipulation, to the extent it occurs, happens on a far smaller scale: No single attack is likely to affect a large number of votes. To make the comparison more explicit, a single teenager, hacker, or other malicious party could potentially affect tens or hundreds of thousands of votes cast through SERVE, while it is extremely unlikely that any single person could conduct vote fraud on such a large scale in existing non-electronic elections. Consequently, vulnerabilities with SERVE could have a far more significant impact than was possible before the introduction of computers and the Internet to the voting process.

**Too Many Potential Attacks.** Because there are many different kinds of attacks that could be conducted against SERVE, as we discuss below, it is essentially impossible to protect against them all. While any particular attack taken in isolation might have a mitigation strategy, the cost of the defense could be high and would be added to the cost of defending against all the other attacks that have been anticipated. Worse yet, defenses created to inhibit one kind of attack may amplify the risks of another. And of course an attack that has not been anticipated remains a serious risk.

**Many Sources of Attacks.** The Internet knows no national boundaries. Consequently, an election held over the Internet is vulnerable to attacks from anywhere in the world. Not only could a political party attempt to manipulate an election by attacking SERVE, but so could individual hackers, criminals, organizations such as the Mafia, and *even other countries*. There is no need to postulate a large conspiracy or highly sophisticated adversaries; many of the attacks we describe could be mounted by lone individuals with college-level training in computer programming.

**Detected vs. Undetected Attacks.** Election fraud has occurred in many different types of elections. A recent example involved boxes of paper ballots that were found floating in San Francisco Bay in November, 2001. There also have been elections in which deceased people have voted. Undoubtedly, many instances of election fraud have gone undetected. But, at least when there is a physical ballot, there is a chance that fraud, or even unintentional mistakes, can be corrected – or at least uncovered.

With Internet voting, however, there is no way to verify that the vote that has been received by SERVE accurately represents the intent of the voter, nor can the voter verify that his/her vote was received by SERVE and accurately recorded by the LEO.

Detected fraud could be almost as worrisome as undetected fraud. If fraud that impacts the SERVE votes were to be detected, it's not obvious what would happen. A judge can order a new election within a jurisdiction, and states can decide how to deal with election related issues. But there is no provision for reholding a federal election. If the 2004 election is as close as the 2000 election was, it is possible that SERVE votes might swing the election to one of the candidates. If there were reason to believe that those votes were unreliable or possibly manipulated, this could have an adverse impact on an already cynical public.

**On-screen Electioneering.** The United States has electioneering laws that prohibit any form of campaigning within some distance of a polling place. In California, for example, that distance is 100 feet. Yet there are no laws yet to prevent the worst kinds of electioneering inside a web browser window while someone is voting. For example, an ISP (or browser company, etc.) may be supported by advertising, like AOL, and would thus have the ability to target ads based on the IP address a user is connected to at the moment. These could take the form of popup ads or even ads within the browser window. The problem is that the moment when a voter connects to the SERVE vote server address (or a voter information site) he/she could be bombarded with all sorts of political ads. It is even possible that at least some forms of ads will end up being protected by the First Amendment, and then there will be no escaping them.

Table 1 describes the threats in terms of the skill level required to mount attacks, the consequences of successful attacks, whether or not the attacks are realistic, and finally the countermeasures that might be used to thwart them.

## **2. Lack of Control of the Voting Environment**

Perhaps the greatest challenge with Internet voting arises from the fact that, in contrast to conventional elections, electoral authorities no longer have control over all the equipment used by voters. With SERVE's Internet voting system, voters can vote from home or elsewhere on their own computers or vote from other locations using computers controlled by third parties. As a result, hackers and other third parties might be able to gain control of a large number of computers used for voting, and election officials would be powerless to protect the integrity of the election. This facet of SERVE's Internet voting architecture poses significant risks to the security of elections. The lack of control of the voting machinery opens up three classes of attack: Compromise of the privacy of votes, disenfranchisement, and vote alteration. The next two sections describe how an attacker could gain control of the voting environment, and what he or she could do once that control is gained.

### ***2.1 How an attacker could control the voting environment***

There are two basic scenarios in which an attacker could control the voting environment: When a voter uses someone else's computer and when a voter's own computer contains malicious software. The latter could occur either because of pre-installed applications designed to attack the election, or because of malicious remote code, such as a worm or virus designed to exploit flaws in the Windows operating system or other applications.

If one votes at a cybercafe, the owners or system administrators of the cybercafe control the computer. In addition, a prior visitor to the cybercafe could have taken control of the computer and installed remote spying or subversion software. There are similar risks to voting from any shared computer such as those at public libraries.

Threat	Skill needed	Consequences	Realistic?	Countermeasures
denial of service attack (various kinds)	low	disenfranchisement (possibly selective disenfranchisement)	common on the Internet	no simple tools; requires hours of work by network engineers; launchable from anywhere in the world
Trojan horse attack on PC to prevent voting	low	disenfranchisement	a million ways to make a complex transaction fail.	can mitigate risk with careful control of PC software; reason for failure may never be diagnosed
on-screen electioneering	low	voter annoyance, frustration, distraction, improper influence	trivial with today's web	nothing voter can do to prevent it; requires new law
spoofing of SERVE (various kinds)	low	vote theft, privacy compromise, disenfranchised voters	web spoofing is common and relatively easy	none exist; likely to go undetected; launchable by anyone in the world
client tampering	low	disenfranchisement	one example: change permissions on cookie file. Many other trivial examples	none exist for all possible mechanisms. Too difficult to anticipate all attacks; likely never diagnosed.
insider attack on system servers	low	complete compromise of election	insider attacks are the most common, dangerous, and difficult to detect of all security violations	none within SERVE architecture; voter verified ballots needed, e.g. Appendix C; likely undetected
automated vote buying/selling	medium	disruption of democracy	very realistic, since voter willingly participates	none exist; buyers may be out of reach of U.S. law
coercion	medium	disruption of democracy	harder to deploy than vote buying/selling, but man in the middle attacks make it achievable with average skill	none exist; likely undetected.
SERVE-specific virus	medium or high	vote theft, privacy compromise, disenfranchised voters	some attacks require only experimentation with SERVE; others require leak of SERVE specs or code and resourceful adversary	virus checking software can catch known viruses, but not new ones; likely to go undetected.
Trojan horse attack on PC to change votes or spy on them	high	vote theft, privacy compromise	widely available spyware would be a good starting point	can mitigate risk with careful control of PC software; harder to control at cybercafe, or other institutionally managed networks; attack likely to go undetected

**Table 1** This table describes, for each potential threat to SERVE, what skill is required by the attacker, the consequences of a successful attack, how realistic the attack is, and what countermeasures might be used to thwart the attack.

If one votes at work, the employer controls the computer. A study found that 62% of major US corporations monitor employee's Internet connections, and more than one-third store and review files on employees' computers.<sup>1</sup> While monitoring Internet connections by passively sniffing would not impact the security of SERVE, since the system uses the Secure Socket Layer (SSL) protocol which encrypts traffic,

<sup>1</sup> [http://www.amanet.org/research/pdfs/ems\\_short2001.pdf](http://www.amanet.org/research/pdfs/ems_short2001.pdf)

any monitoring that is based on software running on the employee's computer could be used for malicious purposes. An employer is also in a position to coerce employees who vote at work into voting a certain way.

Software running on a voter's computer also poses risks. Backdoors, placed in software and activated when a user tries to vote, can invisibly monitor or subvert the voting process. The prevalence of so-called *Easter Eggs* in many popular software packages demonstrates that this is a real possibility. (Easter Eggs are cute extras that a software developer adds to the application without authorization, for fun. One well-known example: Microsoft's Excel 97 spreadsheet application contains a full-fledged flight simulator that can be launched using a secret sequence of keystrokes.)

Today's computers come loaded with software developed by many different entities; any employee at any of those companies could conceivably leave a backdoor that attacks SERVE. Operating systems, games, productivity applications, device drivers, multimedia applications, browser plugins, screen savers, MS office macros are all possible carriers. Every time someone downloads new software, the risk is increased.

In addition to the threat from pre-installed applications, there is a threat from remote attackers. Such an attacker might gain control of a computer without being detected. For example, an attacker could exploit a security vulnerability in the software on a voter's computer. The attacker could then take remote control of the machine using any number of products. Examples of remote control software are PCAnywhere and BackOrifice. It is an unavoidable fact that today's computing systems are inadequate to protect against this threat. Successful penetration of even well-defended computers is routine and common. For example, in December, 2003, the systems at Votehere Inc. were compromised by an external hacker.<sup>2</sup>

Voters' computers in their homes are unlikely to be as carefully defended as corporate ones, and hence voters' machines are especially susceptible to attack. Attacks can be easily automated; hackers routinely scan thousands or even millions of computers in search of those that are easiest to compromise. We can envision scenarios in which the computers of SERVE voters have been compromised on a large scale, calling into question all votes cast over the Internet. Regrettably, such a scenario is all too possible. For example, the recent Slammer worm brought down many ATM machines and compromised many Internet hosts.<sup>3</sup>

Remote attacks might be spread using any of a number of attack vectors. Perhaps most fearsome is a virus or worm that spreads itself and that contains a malicious payload designed to take control of machines and wreak havoc with a future election. Malicious worms are widespread among Internet-connected computers today. For instance, in 2001, the Code Red worm infected 360,000 computers in 14 hours. Modern worms are even more virulent, are often spread by multiple methods, and are able to bypass firewalls and other defenses. Virus checking software programs defend only against previously known viruses, and so are often unable to keep up with the spread of new viruses and worms.

The threat of SERVE-specific viruses should not be discounted. The first question that comes to mind is: "Can virus checking software prevent this threat?" The answer, we believe, is "No." Modern virus checkers are very limited in what they can achieve, because they only attempt to detect previously known viruses. New viruses will not be detected by most current virus checking software. Moreover, it is not too difficult for attackers to build new viruses, or to modify existing viruses sufficiently that they will avoid detection. One can even find virus construction kits on the Internet. In addition, the attacker has the advantage that he/she can test new versions of viruses using the publicly available virus checkers that potential victims use to confirm that the virus will not be detected before its release. In our experience, new viruses usually spread rapidly until their signature is known and the anti-virus companies update their definition files, but this may be too late: the damage to the election may already have been done.

The worm threat is also quite real. It is easy for any competent programmer to write a crude worm; the source code to previous worms can be obtained and modified to create new worms. Writing a sophisticated worm is substantially harder. One set of experts estimates that a small team of experienced programmers could, after months of work, develop a worm that might compromise the majority of all Internet-connected

---

<sup>2</sup> <http://www.cnn.com/2003/TECH/biztech/12/29/voting.hack.ap/>

<sup>3</sup> <http://securityresponse.symantec.com/avcenter/venc/data/w32.sqlexp.worm.html>

computers within a few hours [SPW02]. We don't know if such an ambitious project would succeed on the first attempt, and there seems to be no clear consensus within the security community on how long such a worst-case worm could remain undetected. Some argue that it would be detected within hours or days, while others argue that it may be possible to conceal its existence for weeks or even longer. In any case, worms remain a significant risk. A smaller-scale worm that more selectively targeted a smaller population would be much harder to detect, and possibly could evade detection indefinitely. Even an unsuccessful widespread attack could destroy voter confidence.

There are other ways that attacks might spread. One possible attack involves scanning a large number of computers and attacking them directly. This technique is widely used by hackers. We might also see attacks that are spread by the inclusion of malicious worms in email in order to influence a SERVE election. Elections also could be undermined through the use of web sites that are altered to contain malicious content. Any user visiting such a web site would have his or her computer invisibly hijacked. Because the SERVE voting system requires that voters enable certain dangerous features on their computers, the risk of web site attacks may be heightened. For instance, SERVE supports only Microsoft Windows, a platform that has suffered from many security problems; also, SERVE requires that voters enable ActiveX scripting, cookies, Java, and Javascript—web technologies that pose significant risk to the security of voters' computers.

ActiveX scripting is a Microsoft technology that allows code from the Internet to run natively on client machines. There is a security architecture for ActiveX that is outside the scope of this report. As stated earlier, SERVE requires ActiveX because some of the functionality needed in the system cannot be achieved inside of a browser. However, the use of ActiveX introduces additional vulnerabilities into the system, as shown below.

A dangerous hybrid attack involves placing malicious content on specially chosen websites. For instance, an attacker with a vendetta against one candidate might booby-trap the website of that candidate, so that those who visit the candidate's website are unable to vote using SERVE. Such selective disenfranchisement might eliminate several hundred votes for a candidate, thereby throwing the election to his/her opponent.

There are several ways to booby-trap a website or email that do not pose any major technical difficulties to the attacker. One simple method is to place on the website or in the email a malicious ActiveX control that, when viewed, changes the voter's machine so that it will not work with the SERVE voting system. (We give some examples in the next section.) For a malicious ActiveX control to run, it must be marked as *trusted*.<sup>4</sup> Any programmer who becomes a valid publisher, i.e., whose public key is signed by Microsoft, Verisign, GTE, Thawte or a corporate signing authority, can produce code that is implicitly trusted by the Windows operating system. There have been documented cases of people tricking Microsoft into signing a malicious ActiveX control.

Targeted attacks could be applied either on a large scale or on a small scale. There could be email-driven or web-driven attacks that affect hundreds of thousands or even millions of users: the extraordinary prevalence of spam is testament to the leverage available through email-driven attacks. For instance, these techniques might allow a Republican partisan to prevent a large fraction of Democrats from voting, while leaving most Republicans unaffected – or vice versa. Although it may be possible to build sophisticated email- or web-driven attacks that evade detection, it seems likely that such stratagems would be noticed if employed on a large scale. Nonetheless, even if such an attack were to be detected, there may be little one could do beyond invalidating the entire election, hardly a desirable outcome.

## **2.2 What the attacker can do with control of the voting environment**

An attacker with control over the voter's computer is in position to observe how someone votes, compromising the secrecy of the vote. For instance, such an attacker might place spy software on the computer to silently record all actions taken by the voter. Today, spy software is readily available on the commercial market to record all keystrokes typed, websites visited, and actions taken by the user. What is

---

<sup>4</sup> Code is marked as *trusted* by applying a digital signature with a private key whose public counterpart is part of the Windows system, or whose public component has been certified by Microsoft.

surprising is not that such software exists, but that it is readily affordable to all; one can find such software for under \$50, as well as dozens of free versions. Moreover, if an attacker wants to monitor many voters at once, the software could be customized without too much difficulty to target SERVE elections specifically. Worse, the average computer user would have no way to detect whether third parties have observed his/her vote.

An attacker with control of the voting environment could disable ActiveX or Web cookies, for example, by changing the permissions on the cookie file to disallow write access, so that the user could no longer vote through SERVE. Such an attack is easy to mount. A clever attack could be designed in such a way that the user could not re-enable the necessary permissions. In this case the voter would realize that he/she had been disenfranchised. More sophisticated attacks might cause disenfranchisement in a way that the average voter would not detect.

Targeted voter disenfranchisement poses a serious threat to the integrity of the election. It is possible to imagine widespread attacks that targeted all voters in a particular party for disenfranchisement, leaving the other party unaffected. Such an attack would have serious consequences.

While the ease of selective disenfranchisement is a serious concern, another risk is that a malicious third party with control of a registered voter's computer could use that control to cast an unauthorized ballot, thereby violating the integrity of the election. The fraudulent ballot could appear to come from the authorized voter but in reality be filled out by the attacker. Or, an attacker could wait to see how the voter votes and then change the ballot cast by the voter before it is processed by the ActiveX control from SERVE. Such an attack would require some sophistication, but is not impossible. There are several safeguards in the SERVE design that would hinder this attack. But once the attacker has control of the client, these safeguards could be overcome. The easiest technique might be to modify the ActiveX control running on the machine so that the SERVE safeguards are intercepted and not seen by the user. The *patched* ActiveX control would act as a man in the middle (see below), giving the voter the experience he/she expects when voting legitimately, but modifying the vote. In one possible scenario, the attacker might allow the vote to proceed unmodified if it is to the attacker's liking, modifying or discarding it otherwise.

Privacy compromise, disenfranchisement, and vote fraud potentially could be perpetrated without anyone noticing. The voter might not be aware that his/her vote has been monitored or subverted by a malicious third party. Likewise, election officials would likely not have any way to detect the behavior of such malicious third parties.

### 3. Spoofing and Man in the Middle Attacks

Despite all of the safeguards incorporated in the SERVE system, there are certain attacks that cannot be prevented. In this section, we describe the vulnerability of SERVE to the compromise of voter privacy and ultimately to the subversion of the election. We describe the attacks in increasing order of severity.

The first attack we describe is a man in the middle attack that compromises voter privacy. A man in the middle attack is one in which the adversary interposes itself between the legitimate communicating parties and simulates each party to the other party. To simplify the discussion, we focus on vote privacy, where an attacker seeks to learn how various people voted. The ability of an arbitrary outsider to learn on a wide scale how voters voted is enough of a threat to democracy that we think this alone justifies canceling the SERVE project. The fact that the attack is relatively easy to mount only strengthens our claim.

There are several ways that an adversary could become a *man in the middle*:

- **Control the client machine:** As described in the previous section, if an adversary can control the voting machine, then the adversary can act as a man in the middle to control the vote, even on an encrypted session.
- **Control the local network:** If the attacker has control over the local network environment, such as an employer in a workplace or anyone who shares a wireless network, then the attacker can interpose him/herself as a man in the middle of any network communications.
- **Control an upstream network:** An ISP or foreign government that controls network access from

- the voter to the voting server could masquerade as a man in the middle.
- **Spoof the voting server:** Even without physical access to the network path between a voter and the voting server and without access to the machine, there are social engineering attacks where a voter can be fooled into thinking that he or she is communicating with the voting server. This could be done, for example, by posting or emailing a link that appears to go to the voting server, but in fact does not.
- **Attack the Domain Name Service (DNS):** Attacks against the DNS could route traffic to an attacker instead of to the legitimate vote service.

Clearly, there are many ways that an attacker could become a man in the middle. The use of SSL does little to mitigate the man in the middle attack if the only goal is to compromise privacy. Any man in the middle could act as an SSL gateway, forwarding application data between the voter and the vote server unaltered. However, the attacker would be able to see all of the traffic by decrypting and reencrypting it as it passes between the two. In effect, the attacker would communicate using two SSL sessions, one between itself and the voter, and the other between itself and the vote server, and neither would know that there was a problem.

One attempt to prevent a man in the middle attack would be for the ActiveX control from the voting server to sign the IP address of its SSL endpoint along with the completed ballot. However, this is not a good defense: It would be easy for a man in the middle to defeat this countermeasure by applying a simple patch to the ActiveX control as it traveled to the voter from the server. The patch would hard code the correct IP address into the right spot for the signature. Of course, the attacker would also have to resign the patched ActiveX control; becoming an ActiveX signer requires fooling one of the certifying authorities, or simply purchasing a key.

We carefully analyzed the SERVE system against this attack and conclude that the attack would be relatively straightforward to mount, and that it could be successful.

In the process of analyzing privacy, we discovered another vulnerability of SERVE, where the system could be used for vote selling as follows. The vote seller sets up a proxy server, as described above, and draws voters to his site. In this case the voters willingly connect to the attacker, with the result that the voter's privacy could be compromised. Since the attacker could see how people voted, the vote selling scheme could be verified. Similarly, attackers intent on coercing voters could use the same techniques to learn how the victims had voted, increasing the capability to coerce voters.

Man-in-the-middle attacks could also be used to disenfranchise voters, further raising the level of seriousness of this vulnerability. Once an attacker can act as a man in the middle, the attacker can completely shut out the vote server and spoof the entire interaction with the voter. Although SERVE has some safeguards in place, they assume that the voter knows exactly what to expect from the voting experience; it is probably safe to assume that an attacker could create a voting experience that the voter would believe was real. Once the attacker can spoof the election service, voters are completely disenfranchised. The attacker could make them think that they voted, and the voters will not know that their communications never reached the vote server. One safeguard in SERVE is that voters can check to see if their votes were recorded. It is not clear what the procedure would be if after the election, a large percentage of absentee voters said that they had voted but did not see their names. More likely is the possibility that few voters would bother to check. In either case, the election would be greatly disrupted.

Similar attacks could work against the registration process. Voters could be led to believe that they registered successfully, when in fact they were communicating directly with the adversary and not interacting with the legitimate registration server. The voters would discover when attempting to vote that they were not registered, which could be very disruptive.

Perhaps the most serious consequence of man in the middle attacks is that attackers could engage in election fraud by spoofing the voting server and observing how the voter votes. If a vote is to the attacker's liking, an error message is given and the voter is redirected to SERVE's legitimate voting site; in this case, the vote will be counted. If the attacker does not like the vote, then the entire voting session is spoofed; in this case, the user thinks he/she has voted, but in fact the vote was never received by SERVE and will not be counted. For instance, an attacker could arrange that votes for one candidate will be

received and counted by SERVE, while votes for other candidates will never be seen or counted by SERVE. Thus, the attacker could use the privacy compromise described above to actually subvert the outcome of the election.

While the designers of SERVE are very talented and while they attempted to mitigate against many of the threats posed by Internet voting, man in the middle and spoofing attacks remain threats that are not overcome in the system. It is not clear to us how one would avoid such threats in the current Internet environment. This forms part of the basis for our conclusion that Internet voting cannot be made secure for use in real elections for the foreseeable future.

## 4. Denial of service attacks

If a hacker could overload the election webserver and prevent citizens from voting, the integrity and meaningfulness of the election would be compromised. Such attacks, where legitimate users are prevented from using the system by malicious activity, are known as *denial of service attacks*. We believe that denial of service attacks are a serious risk for SERVE.

Denial of service attacks are possible in everyday life. For instance, flooding a victim's telephone number with a deluge of unwanted phone calls can pose enough of a disruption that the victim disconnects his/her telephone—thereby becoming unreachable to legitimate callers. On the Internet, denial of service attacks are often much more devastating, because Internet denial of service attacks can be automated with a computer, and because such attacks can often be mounted untraceably over the Internet.

A particularly nasty variant of denial of service attack is the *distributed denial of service* (DDoS) attack. In a DDoS attack, many attacking machines collaborate to mount a joint attack on the target. In this scenario, an attacker could take control of many computers in advance by spreading a custom-crafted virus or worm. In computer security jargon, the compromised machines are often known as “zombies” or “slaves,” because the attacker leaves behind hidden software that causes infected machines to blindly obey subsequent commands from the attacker. “Zombie networks” are widely used by hackers today to mount denial of service attacks (and to send spam).

Denial of service is not a theoretical threat; the risk is all too real. Denial of service attacks have become a steadily growing nuisance over the past several years. Automated tools for mounting DDoS attacks have been circulating among the hacker community since at least 1999 [HW01], and hackers routinely amass large “zombie networks” of compromised machines. In February 2000, major DDoS attacks were mounted against several high-profile web sites, including CNN, Yahoo and Ebay.<sup>5</sup> It was later discovered that these damaging attacks had been perpetrated by a lone teenager not on US soil.<sup>6</sup>

Since then, DDoS attacks have become routine. One study recorded over 10,000 denial-of-service attacks during a three-week period in 2001 [MVS01]. In 2001, the Code Red worm infected 360,000 computers in 14 hours; it contained code to mount a DDoS attack on the White House website. (Fortunately, the DDoS attack was deflected at the last minute).<sup>7</sup> In 2003, an Internet election in Canada was disrupted by a denial of service attack on Election Day.<sup>8</sup> These are not isolated examples; it is all too easy to mount DDoS attacks, and the culprits are rarely caught.

### 4.1 How an attacker could mount a denial-of-service attack

Broadly speaking, there are two major forms that Internet denial of service attacks can take. In the first category are attacks in which an adversary is able to swamp the network connection of a targeted webserver with junk data that clogs up the network and prevents other, legitimate traffic from getting through. The second category includes attacks in which the adversary is able to overload the webserver's computational resources with useless tasks that keep it busy; if the webserver becomes too busy, it may be unable to respond to connections from legitimate users. We will discuss both categories.

---

<sup>5</sup> <http://www.nipc.gov/investigations/mafiaboy.htm>

<sup>6</sup> <http://www.cnn.com/2000/TECH/computing/04/18/hacker.arrest.01/>

<sup>7</sup> <http://www.symantec.com/avcenter/venc/data/codered.worm.html>

<sup>8</sup> [http://cbc.ca/stories/2003/01/25/ndp\\_delay030125](http://cbc.ca/stories/2003/01/25/ndp_delay030125)



In a network flooding attack, the adversary sends huge volumes of data in the direction of the victim, saturating the victim's network connection and making it impossible for legitimate users to contact the victim. SERVE's websites are at risk for this kind of attack; if their network connection is overwhelmed by a denial-of-service attack, then eligible voters will not be able to vote using SERVE.

As a rule of thumb, the robustness of a website against network flooding attacks is determined largely by the network capacity available to that website. For instance, a website with a 1 Gbps link to the Internet would be hard-pressed to survive a 1 Gbps DDoS attack. Large e-commerce sites typically have a 10 Gbps link at most. In comparison, researchers have observed DDoS attacks with peak traffic rates in excess of 150 Gbps [MVS01,M+03]. It seems unlikely to us that SERVE could survive such a high-volume DDoS attack.

In the second category of denial-of-service attack, the adversary sends many valid-looking requests to the victim in an attempt to overload the victim's computer and keep it busy with useless work. There are many opportunities for attacks of this sort, and it would be hard to anticipate them all. Instead, we will describe an example of one attack in this category, to give the general idea. Similar principles may apply to many other aspects of the SERVE architecture.

SERVE uses an SSL-protected website. However, SSL is susceptible to a denial-of-service attack. An adversary could send many requests to initiate new SSL connections, and the SSL protocol requires the recipient to perform a slow cryptographic operation (typically an RSA private-key computation) when responding to each such requests. The exact performance depends on the security level provided, but with the fastest and lowest-security key sizes that are today considered acceptable (i.e., 1024-bit RSA), modern computers can handle about 100 new connections per second; hardware accelerators raise this number to thousands of new connections per second. Today's largest e-commerce sites can handle up to 15,000 new SSL connections per second. In comparison, an attacker might be able to initiate about 500,000 new SSL connections per second, based on the following assumption: It is plausible that an attacker could gather a "zombie network" of 10,000 slave computers, and each computer could initiate about 50 new SSL connections per second. Consequently, an attacker could generate 10 to 100 times more SSL traffic than the SERVE website is likely to be able to handle. Thus, a DDoS attack against SERVE's SSL webservers could render SERVE unreachable to voters and disrupt an election in progress.

Unfortunately, mitigating or responding to denial-of-service attacks is very difficult. Today's technology is not up to the task. For example, no good defenses against network flooding attacks are known on today's Internet. It may be possible to defend against the particular SSL attack we describe; however, defending against all variants of this scheme is difficult. As an attacker will attack the weakest link in any system, SERVE must protect against all possible denial-of-service attacks—a very difficult task.

In summary, we are concerned that, no matter how much energy is invested into defensive countermeasures, adequate protection against denial-of-service attacks is unattainable with the technology available today. No matter how careful the designers are, SERVE is unavoidably at risk.

## ***4.2 The implications of denial-of-service attacks for SERVE***

An attacker could mount a large-scale denial of service attack that renders SERVE's voting service unavailable on the day of an election. Those voting on Election Day would be unable to vote, calling into question the validity of the election.

Alternatively, network services could be knocked out or degraded for areas where a particular demographic is known to vote for a particular party. The outcome of the election could be swayed by such an attack. Detection of such an attack would be possible, but it is not clear how to respond; once polls close, there may be no good choices.

SERVE differs from in-person voting in one important aspect: eligible voters can vote at any time during a 30-day window starting 30 days before Election Day and extending until the close of polls on Election Day. If voters could be persuaded to vote early in this time window, the impact of denial-of-service attacks might be reduced: In the past, most denial-of-service attacks have lasted only for a few days, and when the attack subsides, affected voters could then vote, if the polls had not yet closed.

However, we are unconvinced that SERVE's 30-day window can be relied upon to defend against denial-of-service attacks. There are reasons to believe that a large proportion of the voting population will want to vote on Election Day. (See Appendix B for an example). This introduces the threat of *last-day denial-of-service attacks* in which the attacker mounts a denial-of-service attack starting on the morning of Election Day and lasting until polls close. Because responding to denial-of-service attacks takes time, it is likely that an attacker would be able to maintain a last-day denial-of-service attack all day long, so that the SERVE systems would remain unreachable all Election Day. In such a scenario, any overseas citizen who had intended to vote on Election Day would be unable to vote through SERVE, probably would not be able to find any alternative way to vote before the close of polls, and thus would be disenfranchised.

We expect that last-day denial-of-service attacks would disenfranchise a substantial fraction of the SERVE population. There seems to be little that SERVE can do to defend against such attacks. For these reasons, we consider last-day denial-of-service attacks a significant threat to the security of SERVE's elections.

## Conclusions

Our conclusions, based on the arguments in this report are summarized as follows:

- a. The SERVE system has numerous fundamental security problems that leave it vulnerable to a variety of well-known cyber attacks, any one of which could be catastrophic.
- b. Such attacks could occur on a large scale, and could be launched by anyone from a disaffected lone individual to a well-financed enemy agency outside the reach of U.S. law. These attacks could result in large-scale, selective voter disenfranchisement, or privacy violation, or vote buying and selling, or vote switching even to the extent of reversing the outcome of many elections at once, including the presidential election. With care in the design of the attacks, some of them could succeed and yet go completely undetected. Even if detected and neutralized, such attacks could have a devastating effect on public confidence in elections.
- c. It is impossible to estimate the probability of a successful cyber-attack (or multiple successful attacks) on any one election. But we show that the attacks we are most concerned about are quite easy to perpetrate. In some cases there are kits readily available on the Internet that could be modified or used directly for attacking an election. And we must consider the obvious fact that a U.S. general election offers one of the most tempting targets for cyber-attack in the history of the Internet, whether attacker's motive is overtly political or simply self-aggrandizement.
- d. The vulnerabilities we describe cannot be fixed by design changes or bug fixes to SERVE.. These vulnerabilities are fundamental in the architecture of the Internet and of the PC hardware and software that is ubiquitous today. They cannot be eliminated for the foreseeable future; indeed, if they were to be eliminated, it would be only as a result of wholesale redesign and replacement of much of the hardware and software security systems that are part of, or connected to, today's Internet.
- e. We have examined numerous variations on SERVE in an attempt to recommend an alternative system that may deliver somewhat less voter convenience in exchange for fewer or milder security vulnerabilities. However, all such variations suffer from the same kinds of fundamental vulnerabilities that SERVE does; regrettably, we cannot recommend any of them.
- f. The SERVE system might appear to work flawlessly in 2004, with no successful attacks detected. Unfortunately, a seemingly successful voting experiment in a U.S. presidential election involving seven states would be viewed by most people as strong evidence that SERVE is a reliable, robust, and secure voting system. Such an outcome would encourage expansion of the program by FVAP in future elections, or the marketing of the same voting system by vendors to jurisdictions all over the United States, and other countries as well.

However, the fact that no successful attack is detected does not mean that none occurred. Many attacks, especially if cleverly hidden, would be extremely difficult to detect, even in cases when they change the outcome of a major election. Furthermore, the lack of a successful attack in 2004 does not mean that successful attacks would be less likely to happen in the future; quite the

contrary, future attacks would be more likely, both because there is more time to prepare the attack, and because expanded use of SERVE or similar systems would make the prize more valuable. In other words, a “successful” trial of SERVE in 2004 is the top of a slippery slope toward even more vulnerable systems in the future.

- g. In spite of the fact that we recognize the need to offer better support for voting for overseas Americans, we are forced to conclude that the best course to take is not to field the SERVE system at all. Because the danger of successful, large-scale attacks is so great, we reluctantly recommend shutting down the development of SERVE and not attempting anything like it in the future until both the Internet and the world’s home computer infrastructure have been fundamentally redesigned. In Appendix C of this report, we offer an alternative to SERVE that we believe is much more secure.

We want to make clear that in recommending that SERVE be shut down, we mean no criticism of the FVAP, or of Accenture, or any of its personnel or subcontractors. They have been completely aware all along of the security problems we described, and we have been impressed with the engineering sophistication and skill they have devoted to attempts to ameliorate or eliminate daunting security problems. We do not believe that a differently constituted project could do any better job than the current team. The real barrier to success is not a lack of vision, skill, resources, or dedication; it is the fact that, given the current Internet and PC security technology, the FVAP has taken on an essentially impossible task. There really is no good solution. The project is thus too far ahead of its time, and should wait until there is a much improved security infrastructure to build upon.

## Acknowledgements

We thank Doug Jones, Deirdre Mulligan, and Ron Rivest for helpful comments.

## References

- [AK96] Ross Anderson and Markus Kuhn, “Tamper Resistance - a Cautionary Note,” *Proceedings of the Second Usenix Workshop on Electronic Commerce*, pages 1-11, November, 1996.
- [HW01] Kevin J. Houle, George M. Weaver, “Trends in Denial of Service Attack Technology”, October 2001.
- [Khono03] Tadayoshi Khono, Aviel D. Rubin, Adam Stubblefield, and Dan S. Wallach, “Analysis of an Electronic Voting Machine”, Johns Hopkins Information Security Institute Technical Report TR-2003-19, July 23, 2003.
- [MVS01] David Moore, Geoffrey M. Voelker, Stefan Savage, “Inferring Internet Denial-of-Service Activity”, *Usenix Security* 2001.
- [M+03] David Moore, Vern Paxson, Stefan Savage, Colleen Shannon, Stuart Staniford, Nicholas Weaver, “Inside the Slammer Worm”, *IEEE Security & Privacy* 2003.
- [Pfleeger03] Charles P. Pfleeger and Shari Lawrence Pfleeger, *Security in Computing third edition*, Prentice Hall, 2003.
- [SPW02] Stuart Staniford, Vern Paxson, Nicholas Weaver, “How to Own the Internet in Your Spare Time”, *Usenix Security* 2002.

## Appendix A

The main body of the report covers some of the most serious security risks in great detail, but we did not attempt to give an exhaustive list. This appendix is intended to supplement that list. In this appendix, we briefly discuss several other security issues that also pose serious risks for SERVE.

### ***Vulnerabilities in Servers Could Breach Election Security***

SERVE uses centralized computers (servers) to record and forward votes. If those servers are compromised, every ballot cast through SERVE could be modified or replaced, and the integrity of the entire election would be irreparably damaged. Since the servers are a central single point of failure, it is absolutely vital that they resist attack.

The risk of intrusion into SERVE's centralized computers is, unfortunately, significant. SERVE has deployed a careful and well-designed firewalling architecture designed to prevent many kinds of direct attacks; however, there remain possible vulnerabilities in the software exposed to the outside world that could enable attackers anywhere on the Internet to penetrate SERVE's defenses and gain control of the servers. Some examples of the kinds of vulnerabilities are buffer overruns, format string vulnerabilities, directory traversal bugs, race conditions, cross-site scripting bugs, SQL injection bugs, cryptographic failures, and session authentication weaknesses. It is important to understand that there are no comprehensive methods for avoiding, or testing for, these vulnerabilities and bugs. As a result, even the most widely used, exhaustively-tested software in the world usually has many such vulnerabilities. SERVE has deployed several mitigation strategies that we will not enumerate here, but these strategies are not sufficient, in our view, to reduce the risk to acceptable levels.

High-confidence software can be found in safety-critical systems such as nuclear reactor control subsystems, fly-by-wire passenger airplanes, the FAA's air-traffic control system, the Space Shuttle, and certain military applications. Designers of safety-critical systems typically avoid the use of commercial software, because it is widely accepted that standard commercial programming practices pose an unacceptable risk for such applications. Designers of safety-critical software employ known techniques for building highly reliable software. These elaborate and costly techniques have not been used in the development of SERVE. Therefore, we must recognize that the standard commercial practices used in SERVE, and in the commercial-off-the-shelf (COTS) software it builds upon, come with an unavoidable risk of failure.

### ***Existing Processes are Inadequate for Certifying Voting Software***

Election law in most states requires that all voting systems—whether electronic or not—be qualified by an authorized federally licensed laboratory known as an Independent Testing Authority (ITA), and then submitted to the state for certification. The ostensible purpose of these procedures is to make sure that the voting system meets the voluntary federal voting system standards promulgated by the FEC (and in the future, by NIST), and that they conform to the state's election laws. It is tempting to place a lot of faith in certification procedures as a means for preventing security failures. We believe such faith is unwarranted. We argue that even a lengthy, conscientious testing and examination program by the most qualified people cannot give us the necessary security guarantees. In fact, in general, no process can, since in most cases the problem of establishing that a program meets any particular security requirement is known to be fundamentally unsolvable [Pfleeger03].

With electronic and Internet voting systems the most important part of the ITA qualification process is review and testing of the software. We argue however, that this review and testing does not, and cannot, offer much guarantee at all that the software in voting systems actually does the job it is supposed to do. Ideally one would like the testing to be designed to verify that the software is at least *correct*, i.e. that it captures and counts votes properly under all normal conditions, in any possible election, and with any legal voter behavior. It should also verify that the software is *robust* and *reliable*, i.e. that it works reasonably even in the presence of various bug and failure scenarios, including aberrant behavior by users (such as clerks and voters, etc.); and that the software is *secure*, i.e. that it contains no internal malicious logic (Trojan horse code) and is not vulnerable to any of a vast range of potential external attack scenarios.

Unfortunately, however, neither the ITAs nor anyone else could perform review and testing even remotely comprehensive enough to establish any of these properties because the time and labor involved would be truly astronomical. There are fundamental limits to what testing can accomplish; it is a commonplace of the software world that *while testing can be used to verify that bugs and security vulnerabilities are present, it can never prove that they are absent.*

What the ITAs actually do to qualify voting system software is considerably less than that. They do run mock elections with test votes to verify that the software appears to work as required under normal conditions; this amount of testing is probably sufficient to detect simple program errors and obvious defects, although those will likely already have been found by the developers. (We do not really know because, as we note later, results of this testing in SERVE are secret.) But generally speaking such a testing process cannot be expected to do more than filter out the most obvious bugs, leaving intact more subtle bugs that get triggered less frequently. And it is safe to say that the ITAs do very little testing at all to look for software reliability, security, or malicious code problems, since testing is generally ineffective for those purposes.

Besides testing, the ITAs also examine the source code of the software. Most of the examination is in the form of automated scripts run over the source that look for simple properties required by the FEC standards, e.g. that there are sufficient comments in the code, that “modules” in the programs are not too long, and that each has a single entrance and exit. These are syntactic and stylistic requirements that are only crude indicators of good engineering practices; they indicate nothing at all about the code’s correctness, reliability, or security.

We have been told that the ITA engineers also examine the source code by eye, searching for suspicious-looking code fragments that might indicate the presence of a bug, a security vulnerability, or malicious logic. If so, it may give ITA engineers a better idea of the software engineering skill that went into the program development. But it is unlikely that someone not on the development team will spot many subtle bugs in a large codebase because he or she can never take the time to fully understand its overall code structure in the necessary detail. And, contrary to many people’s intuition, it is unlikely in the extreme that anyone, whether on the development team or not, would detect malicious logic that was deliberately disguised by a clever programmer, no matter how much effort was put into the search.<sup>9</sup> It is much easier to hide a needle in a haystack than to find it.

However, even if the ITAs were to do an excellent job of examining the software used in election systems, we would still have major concerns. Unfortunately, the ITAs operate under the same cloak of secrecy as voting machine vendors. Both the tests conducted by the ITAs and the results of those tests are secret. Still, there are some things that we know about how ITAs have evaluated computerized touch screen machines (DREs), and there are other things that we can deduce from some glaring failures that have been exposed in the DREs. It is reasonable to assume that the shortcomings of ITAs with respect to DREs will carry over to their certification of Internet voting.

For example, there are fundamental differences between modern software systems and previous voting systems. In particular, today’s software-based systems are orders of magnitude more complex than mechanical and paper voting machines, and this complexity poses tremendous challenges for verification. In a mechanical voting machine, there are only so many moving parts and only so many ways that the system can misbehave. In contrast, modern software systems are exceptionally complex, with the electronic equivalent of hundreds of thousands or millions of moving parts; such machines can fail in convoluted and unpredictable ways. The very fact that Microsoft and other software vendors are forced to issue frequent software patches (see below) demonstrates that the complexity of modern software systems makes it impossible to detect all software errors, let alone malicious code, with testing and code reviews. Because current methods are good at exposing only simple failures, these methods are not a reliable way of certifying software.

Even if current practices were effective at checking that the voting system behaves as it should during normal operation, security is by definition about the behavior of the system when it is under attack by a malicious entity, a (hopefully) abnormal situation. Because it is hard to anticipate how attackers might

---

<sup>9</sup> For an example of software problems that were undetected by ITA testing, see [Khono03].

behave, it is hard to test for security flaws. Existing practices can almost certainly be improved, but there are limitations on what can be achieved given our current state of knowledge.

It is even harder to detect intentional security flaws that have been purposefully planted by insiders. After careful consideration, we have concluded that if a malicious insider with access to SERVE software wanted to insert a back door in SERVE's software systems, it is likely that the attacker could camouflage the attack well enough to avoid detection by the ITA's during the certification process, especially in the absence of a meaningful code review.

As if these problems with the certification process were not enough, there is a giant loophole in the FEC requirements: commercial-off-the-shelf software does not have to be tested at all by the ITAs. It is simply assumed to be bug-free, Trojan free, and invulnerable to external attack, all of which are manifestly false.

And finally, we note that for systems where security is critical, testing should be done in a truly hostile environment. As Anderson and Kuhn point out in [AK96], "although it is necessary to design commercial security systems with much more care, it is not sufficient. They must also be subjected to hostile testing." The SERVE system has undergone ITA testing but has not existed in a truly hostile environment.

Given these difficulties, it is no surprise that recent experiences with certification of software voting systems have not been encouraging. Regrettably, there are no good solutions in sight. The computer security community has been struggling with these problems for decades, and they are unlikely to be resolved any time soon.

### ***Commercial-Off-The-Shelf Software Poses a Major Risk to Election Security***

SERVE relies heavily on commercial-off-the-shelf (COTS) software. For instance, voters will vote from computers running the Microsoft operating system. SERVE's infrastructure is built on top of commodity operating systems and applications. The reliance on COTS software introduces significant risks.

COTS software is exempted from testing or other evaluation by ITA's. Because most of the software running on SERVE computers is COTS, not custom-designed software, most of the software on the election servers will not be examined or evaluated during the certification process. SERVE machines will run millions lines of code of COTS software, but all this code is exempt from testing. Even the cryptographic core (that is, the software controlling the cryptographic operations) of SERVE is exempt, due to its COTS status.

Much of our modern COTS software is plagued by security problems, and new security vulnerabilities in COTS software are discovered daily. For instance, over 4,000 new COTS vulnerabilities were reported in 2002 alone. Such considerations make it likely that the COTS software that SERVE is relying upon will have unknown vulnerabilities. These vulnerabilities could compromise election security if discovered and exploited.

This raises the question: "What do we do if a new security vulnerability is discovered in software upon which SERVE relies?" Currently, when a security hole is discovered, the software vendor issues a security patch. Such patches are issued on a fairly frequent basis. Since security patches modify the voting system, SERVE is required to test the security patch and to submit the code to the ITA for certification. This procedure is expected to take about two weeks from patch issuance to patch installation.

Today, it is routine to see hackers exploiting new security vulnerabilities only days after they are initially announced or after the patch is first made available. Thus, on the one hand a two-week delay may be too long to prevent exploitation of vulnerabilities. On the other hand, two weeks may be insufficient to install, *adequately* test, and certify a new patch.

### ***Why Internet Voting is far more Difficult than e-Commerce***

Some advocates of commodity software have argued that COTS software is everywhere; it is the most widely used software in the world, and it is the most thoroughly tested software in the world. These advocates suggest that the exemption for COTS software from testing or other certification requirements is

made tolerable by the prevalence of COTS software.

We believe that the pro-COTS argument is fundamentally mistaken. We agree that COTS software is widely used, but we disagree that widespread usage is reason for confidence in its security. Indeed, broad usage of COTS software has only underscored that COTS software and the commodity applications used on the Internet are rife with numerous security defects. Security defects in COTS software may be tolerated in the e-commerce world, but voting has very different requirements than e-commerce.

First, security is essential to elections. Democracy relies on broad confidence in the integrity of our elections, so the stakes are high. We cannot afford to get this wrong. Consequently, voting requires a higher level of security than e-commerce. Though we know how to build electronic commerce systems with acceptable security, *e-commerce grade security is not good enough for elections*.

Second, securing Internet voting is structurally different—and fundamentally more challenging—than securing commerce over the Internet. For instance, it is not a security failure if my spouse uses my credit card with my consent, but it is a security failure if my spouse can vote on my behalf.

Third, the anonymity requirement for voting makes electronic voting a fundamentally more difficult problem than e-commerce. For example, in a commercial setting, customers can detect and deter fraud by checking their bill or receipt, and when fraud is detected, it is possible to recover from this failure. In contrast, voting systems cannot provide receipts, because receipts would violate voter anonymity and would enable vote-buying and coercion attacks.

Fourth, in the e-commerce world, credit card companies effectively act as insurers who absorb the cost of certain types of fraud and make up for these losses by charging merchants transaction fees. In the voting world, there is no obvious way to insure against security failures; how do we recover if an election is compromised?

Internet voting poses demanding requirements on our software, and commodity software that is perfectly acceptable in an e-commerce setting may not be acceptable for Internet elections. Therefore, we believe that the heavy use of COTS software in SERVE, the known security risks of COTS software, and the exemptions from certification for COTS software represent significant risks for SERVE.

## Appendix B

### ***The ICANN experience***

Internet voting has been occurring for several years, most typically with private elections, such as stockholder elections or the election of officers of private organizations.

Internet elections might possibly be acceptable if the stakes are not high. For example, if the issues on which stockholders are voting are not controversial, then there is little incentive to subvert a stockholder election. By contrast, there is a lot of incentive to subvert national presidential and congressional elections.

In 2000 the Internet Corporation for Assigned Names and Numbers (ICANN)<sup>10</sup> held an election over the Internet. Anyone age 16 or older with an Internet address was eligible to vote. The world was divided into five regions, and candidates were nominated from each region by a nominating committee. There was also a Member Nomination process that involved the “endorsement” of a potential candidate via the Internet.

Problems with accessing the official website occurred in every phase of the ICANN election, including voter registration, endorsement, and voting. ICANN significantly underestimated the computing resources that would be required, especially since people tended to wait until near the various deadlines to attempt to access the ICANN website. Also, some of the processes that ICANN had put in place to try to minimize fraud, such as requiring “activation” of membership after a member had registered, were confusing to voters. Many voters were unable to vote because they did not activate their membership by the deadline.

Even though individual passwords were mailed to physical addresses, some people claim to have voted multiple times. There were also numerous accounts of people being disenfranchised. Disenfranchisement occurred because the ICANN website was overwhelmed by the demand, because some voters never received their passwords or subsequently lost them, and because some voters did not realize that registration was not sufficient for voting. In spite of difficulties with registration, about 158,000 people registered. While 76,183 of the registered voters activated their membership, the number who managed to vote in the election was only 34,035.<sup>11</sup> These numbers suggest large-scale disenfranchisement of voters at every stage of the process, with fewer than 1/4 of the initially registered voters actually voting. According to the Markle Foundation report, “the technical weaknesses in the registration system made it virtually impossible to assess the integrity of the voters’ list, the security of the PINs, and secrecy of vote.”

There were also rumors of country or corporate competition to pack the voter list during the later phases of the registration period. Technical problems experienced during the first day of voting, which took place over a ten day period, created, according to the Markle report, “a credibility problem...with voters and interested parties watching the process.”

---

<sup>10</sup> One of the authors of this paper, Barbara Simons, was nominated by the endorsement process; she was the runner-up in North America, losing to Karl Auerbach.

<sup>11</sup> “Report on the Global, On-Line, Direct Elections for Five Seats Representing At-Large Members on the Board of Directors, a report by the Markle Foundation, [http://www.markle.org/News/Icann2\\_Report.Pdf](http://www.markle.org/News/Icann2_Report.Pdf)



## Appendix C

### ***An Alternative to SERVE***

Although we are concerned that the SERVE architecture has too many serious vulnerabilities to be used in public elections, we agree nonetheless that there is a need to lower barriers to voting for overseas Americans and the military. Here we offer an alternative architecture that we believe has most of the benefits that SERVE offers, with far less vulnerability. We are not proposing that this scheme be adopted, but we do suggest that it is a much better starting point for a voting system for that population.

The troubles with SERVE derive from three fundamental design choices: It uses the Internet heavily, with all of the vulnerabilities that implies (e.g. denial of service, spoofing, and man-in-the-middle attacks). It relies on voters using private, unsecured PCs with proprietary, commercial software configured to accept mobile code, with all of the vulnerabilities that implies (e.g. virus attacks, various kinds of privacy violation). And SERVE itself is proprietary software, with all of the vulnerabilities that implies (e.g. security holes, bugs, insider fraud).

We suggest building a system that avoids the dangers of these architectural features. Here is how it might be done:

- a) Base the system on kiosks, to be located at consulates and military bases at major sites around the world. Voters would come to the kiosk to vote, rather than vote from their own PCs.
- b) The software on the kiosk would be booted from a clean copy, maintained and configured by trained elections officials so that the software environment on the voting machine is known and controlled.
- c) The kiosk is never connected to the Internet; hence, no Internet-related attacks are possible. It receives software and databases via disks (such as DVD Write Once Read Many (WORM)) delivered by certified mail in advance of the election. It does not transmit ballots back to the counties by Internet; rather, it prints them, and they are mailed back to the counties, just like any other absentee ballot.
- d) The kiosk has three databases that allow it to perform the voting functions. All three together could fit on one DVD-WORM disk.
  - i) An identification database to authenticate all voters who have registered to vote via the SERVE system. This would have at most 6 million records, since that is the size of the eligible population.
  - ii) A voter registration database to indicate which ballot image (style) each voter is supposed to receive. This data is compiled from information contributed by any or all of the 3000+ county jurisdictions in the U.S. that wish to participate. There are approximately 200 million registered voters in the U.S., but only at most the 6 million records corresponding to those eligible to vote via SERVE would be needed.
  - iii) A ballot image database, containing a PDF or XML representation of each ballot style used in every county in the U.S. that is participating in SERVE. There are on the order of 100,000 such ballot styles. These also would be provided by the LEO in the counties participating in SERVE.

The voter at the kiosk, identifies himself or herself using military ID or other authentication information provided by SERVE. The kiosk verifies that the voter is registered to vote, and looks up which ballot style to issue. Then the voter indicates his or her choices on a touch screen, and prints out a completed ballot. The voter examines the printed ballot to make sure it correctly indicates his or her choices. (If not, the voter approaches the election official for help in voiding the ballot and re-voting.) This last step provides for a *voter-verified paper ballot*, and assures that no potential bugs or Trojan logic anywhere in

the system's code can incorrectly record the voter's intent. Finally, the voter deposits the completed ballot in an envelope addressed to his home county for mailing back, again like a traditional absentee ballot. A system like this is essentially a remote ballot printer located near most Americans anywhere in the world. Its design would have to be fleshed out with more detail and additional security procedures. These include preventing multiple voting, managing encryption keys for the databases, and determining what kinds of electronic record of the voting transactions the machines keep and what they do with these records. Such a system would reduce the voting transaction from three trips through the mail down to one, and is vastly safer than SERVE from any kind of remote or programmed attack.

## Appendix D

### **Author Biographies**

**David Jefferson** is a computer scientist at Lawrence Livermore National Laboratory where he does research in scalable parallel computation (supercomputing). He has also worked for many years in the field of election technology and security, serving as chair of the Technical Committee of the California Secretary of State's Task Force on Internet Voting in 1999-2000; as a member of the executive committee of the NSF panel on Internet voting in 2000-2001; and a member of the California Secretary of State's Task Force on Touchscreen Voting in 2003. He has been active nationally as a speaker and advisor on computerized voting system issues, and is a current member and past chair of the board of directors of the California Voter Foundation. Before joining LLNL, he was a computer scientist at DEC/Compaq Labs in Palo Alto, and before that had been Associate Prof. of Computer Science at UCLA, where he was best known as the co-inventor of the Time Warp method of parallel discrete event simulation.

**Aviel D. Rubin** is Associate Professor of Computer Science and Technical Director of the Information Security Institute at Johns Hopkins University. Prior to joining Johns Hopkins Rubin was a research scientist at AT&T Labs. Rubin is author of several books including *Firewalls and Internet Security*, second edition (with Bill Cheswick and Steve Bellovin, Addison Wesley, 2003), *White-Hat Security Arsenal* (Addison Wesley, 2001), and *Web Security Sourcebook* (with Dan Geer and Marcus Ranum, John Wiley & Sons, 1997). He is Associate Editor of *ACM Transactions on Internet Technology*, Associate Editor of *IEEE Security & Privacy*, and an Advisory Board member of Springer's *Information Security and Cryptography Book Series*. Rubin serves on the board of directors of the USENIX Association and on the DARPA Information Science and Technology Study Group.

**Barbara Simons** Put bio info here

**David Wagner** is an Assistant Professor in the Computer Science Division at the University of California at Berkeley. He has extensive experience in computer security and cryptography, and over 50 technical publications. He and his Berkeley colleagues are known for discovering a wide variety of security vulnerabilities in various cellphone standards, 802.11 wireless networks, and other widely deployed systems. In addition, David was a co-designer of one of the Advanced Encryption Standard finalists, and he remains active in the areas of systems security, cryptography, and privacy. David is an Alfred P. Sloan Research Fellow and a CRA Digital Government Fellow.