

[Go to the U of M home page](#)
[OneStop myU](#)

HUMPHREY SCHOOL OF PUBLIC AFFAIRS



[HOME](#)
[PROGRAMS](#)
[RESEARCH](#)
[TOPICS](#)

Cautionary Tale: Student Gets Jail Time for Stealing Online School Election

By [Doug Chapin](#) on July 17, 2013

[SHARE](#)



[Image courtesy of [black-of-hat](#)]

Technically, this isn't the kind of election news I usually blog about (because it doesn't involve a public election) but I thought it was worth sharing ... From [UTSanDiego](#):

A former Cal State San Marcos student who rigged a campus election by stealing nearly 750 student passwords to cast votes for himself and friends was sentenced Monday in federal court to a year in prison ...

Weaver, 22, of Huntington Beach was a third-year business student when he carried out the elaborate plan to win election as president of the school's student council in March 2012. He pleaded guilty this year to three federal charges, including wire fraud and unauthorized access to a computer ...

The plan to steal the election was months in the making.

On Weaver's computer, authorities found a PowerPoint presentation from early 2012, proposing that he run for campus president and that four of his fraternity brothers run for the four vice president spots in the student government. The presentation noted that the president's job came with an \$8,000 stipend and the vice presidents each got a \$7,000 stipend.

Weaver also had done a bit of research, with computer queries such as "how to rig an election" and "jail time for keylogger."

A month before the election, Weaver purchased three keyloggers -- small electronic devices that secretly record a computer user's keystrokes [pictured above - ed.].

Authorities said Weaver installed keyloggers on 19 school computers, stole passwords from 745 students and cast ballots from the accounts of more than 630 of those victims.

The plot was discovered, however, when technicians spotted unusual activity on the last day of the election period:

Using remote access, technicians watched the computer user cast vote after vote. They also watched as the user logged into the account of a university official and read an email from a student complaining that the system would not let her vote.

Weaver had already cast a ballot from the student's account, which was why she couldn't vote.

The techs called campus police, who found Weaver at the school computer. He had keyloggers with him and was arrested.

The student didn't help himself when he engaged in an elaborate cover-up afterwards:

After a brief jail stay, Weaver and a friend created fake Facebook pages using the names of real students. They posted fictitious conversations on those pages to make it look as if those students had conspired to frame him.

The "conversations" on those bogus pages were sent to reporters at U-T San Diego, 10News and the campus newspaper but none took the bait.

Indeed, it was the cover-up more than the crime that earned Weaver jail time:

"That's the phenomenal misjudgment I can't get around," said Judge Larry Burns, who rejected Weaver's request for probation.

Burns said the election rigging was a serious offense but "kind of juvenile." Developing a scheme to deflect blame after he had been caught made it worse.

"He's on fire for this crime, and then he pours gasoline on it to try to cover it up," the judge said.

Now, as I said above, this story isn't typical of the kind of news I usually cover - but I still think it's appropriate for a few reasons:

- + First, election officials and their staffs need to be aware of the existence of keyloggers even though online voting is not used anywhere in the nation. Specifically, giving an intruder inside access to an election management system leaves the entire process vulnerable to tampering. You and your staffs should be able to identify a keylogger - and if they're in use make sure you know by whom and why;
- + Second, even when an attempt to rig an election is unsuccessful, it still generates clean-up costs for the affected entity. The University incurred \$40,000 in costs to close the security breach in this case - costs that would likely be dwarfed in the event of an attack on a real election.

Stories like this are in many ways a blessing to the field because they highlight potential vulnerabilities; if nothing else, election offices should stay on top of these developments because it's almost a certainty that a potential attacker will.

3 Comments

AndyK

July 17, 2013 10:32 AM | [Permalink](#) | [Reply](#)

I guess this just goes to show that online voting will never work. It's way too risky. My vote won't be counted because someone has already voted for me, or because a million other fake people voted, etc.

Voter fraud should be expensive and slow (see Al Franken). Online voting just makes it even easier.

William J. Kelleher, Ph.D.

July 18, 2013 4:12 PM | [Permalink](#) | [Reply](#)

Hi Doug!

David Jefferson was inspired by your post to spin some scary stories about Internet voting's, in your words, "potential vulnerabilities." His post is at

[http://electionlawblog.org/?](http://electionlawblog.org/?p=53082&utm_source=feedburner&utm_medium=email&utm_campaign=Feed%3A+electionlawblog%2FuqCP+%28Election+Law%29)

[p=53082&utm_source=feedburner&utm_medium=email&utm_campaign=Feed%3A+electionlawblog%2FuqCP+%28Election+Law%29](http://electionlawblog.org/?p=53082&utm_source=feedburner&utm_medium=email&utm_campaign=Feed%3A+electionlawblog%2FuqCP+%28Election+Law%29)

I wrote a reply, Student Election Hack – Is it a Bad Omen for Internet Voting?

at

<http://internetvotingforall.blogspot.com/2013/07/student-election-hack-is-it-bad-omen.html>

Thanks,

Bill Kelleher

[Online voting](#)

August 3, 2013 2:30 AM | [Permalink](#) | [Reply](#)

He's on fire for this misdeed, and then he pours petrol on it to try to cover it up. Thanks for sharing this. It will be helpful for many other people and they will get very informative information about topic.

Leave a comment

Name

Email Address

URL

Remember personal info?

Comments (You may use HTML tags for style)



Type the text

[Privacy & Terms](#)



- [CONNECT](#)
- [ABOUT](#)
- [SEARCH](#)
- [ARCHIVES](#)

 [Subscribe in a reader](#)

 [Subscribe by email](#)

 [Twitter](#)

 [Contact the author](#)

HUMPHREY SITES

**CENTER for the STUDY of
POLITICS and GOVERNANCE**

Election Law Blog

Election Law Blog

The law of politics and the politics of law: election law, campaign finance, legislation, voting rights, initiatives, redistricting, and the Supreme Court nomination process ◊ Rick Hasen's blog

David Jefferson on College Vote Hack

Posted on [July 17, 2013 10:00 am](#) by [Rick Hasen](#)

[David Jefferson](#) sends along these observation:

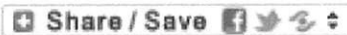




I just read [Doug Chapin's article](#) on the vote rigging at Cal State San Marcos, and I would add several observations. Had this been a public election conducted via Internet voting, it would have been much more difficult to identify any problem or to capture the perpetrator, Mr. Weaver.

Mr. Weaver was captured because he was voting from school-owned computers. This was networked voting but not really Internet voting. The IT staff was able to notice "unusual activity" on those computers, and via remote access they were able to "watch the user cast vote after vote". But in a public online election people would vote from their own private PCs, and through the Internet, not on a network controlled by the IT staff of election officials. There will likely be no "unusual activity" to notice in real time, and no possibility of "remote access" to allow them to monitor activity on a voter's computer. Note also that university IT staff were able to monitor him while he was voting, showing that they were able to completely violate voting privacy, something we cannot tolerate in a public election.

In the Cal State San Marcos election votes apparently had to be cast from computers on the university's own network, and not from just anywhere on the Internet. I infer this because it makes good security sense, and because I cannot think of any other reason Mr. Weaver would cast his phony votes from a university computer rather than from an anonymous place like a public library. If this is correct, it is a huge security advantage not possible in public elections, where the perpetrator could be anywhere in the world. Even if public officials somehow did notice an unusual voting pattern that made them suspicious after the fact that phony votes were cast, there would be no evidence to indicate who it was, and no police on the spot to pick him up red handed.

Even with the restriction to voting on university-managed computers Mr. Weaver could still have gotten away with his crime, but he made several mistakes. He cast hundreds of phony votes one by one, in person, while sitting at the keyboard of the computer. But if he had been a programmer and been a little smarter he would have used one of his captured passwords to log in as someone else, and then (as that person) he could have run a program to automate the casting of all of those phony votes, possibly even from other machines. He could have scheduled the script to run an hour later when he was long gone, and with a random delay between each cast vote so that detection of a suspicious pattern would be less likely. The scheme might still have been caught anyway, but there would have been no technical evidence implicating Mr. Weaver.

When he was arrested Mr. Weaver was caught with very suspicious hardware keylogging devices on his person that he had used to capture people's passwords. Again, if he had been a little smarter, he would have used one of his keylogging devices to capture the password of a system administrator, and then used that password to install keylogger software on other campus computers to capture the students' passwords. Then he would not have had any hardware keyloggers on him and would not have been tied to the crime that way. Or if he had been able to capture a root password he might have been able to use it to simply edit the database of recorded votes, and not bothered with actually casting any phony votes. Finally, we have to note that even though Mr. Weaver was caught, there was apparently no way to know for certain (other than taking his word for it) exactly which votes were the phony ones he cast. They could not simply remove his votes from the count and let the election be certified based on the remaining votes. They had to cancel and re-run the entire election. Needless to say, this would not be an easy option in a public election, which can only be cancelled after the fact by a court after lawsuits and potential appeals. The legal process in such a case would likely be a protracted mess and permanently damage public confidence in elections (again). In the many debates on the subject of Internet voting it is important not to allow anyone to use this Cal State San Marcos student election experience to argue that online public elections can be made safe because those who would cast phony votes will be caught. Mr. Weaver's actions were detected because he was voting from computers controlled by the university IT staff, and he was identified and caught because he was not even minimally technically skilled in the techniques that could have distanced him from the crime. In a high stakes public election we will not be so lucky.

 Share / Save    

This entry was posted in [internet voting](#), [voting technology](#) by [Rick Hasen](#). Bookmark the [permalink \[http://electionlawblog.org/?p=530821\]](http://electionlawblog.org/?p=530821).

This copy is for your personal, noncommercial use only. You can order presentation-ready copies for distribution to your colleagues, clients or customers, please [click here](#) or use the "Reprints" tool that appears next to any article. Visit www.nytreprints.com for samples and additional information. Order a reprint of this article now. »



January 21, 2004

Report Says Internet Voting System Is Too Insecure to Use

By JOHN SCHWARTZ

A new \$22 million system to allow soldiers and other Americans overseas to vote via the Internet is inherently insecure and should be abandoned, according to members of a panel of computer security experts asked by the government to review the program.

The system, Secure Electronic Registration and Voting Experiment, or SERVE, was developed with financing from the Department of Defense and will first be used in this year's primaries and general election.

The authors of the new report noted that computer security experts had already voiced increasingly strong warnings about the reliability of electronic voting systems, but they said the new voting program, which allows people overseas to vote from their personal computers over the Internet, raised the ante on such systems' risks.

The system, they wrote, "has numerous other fundamental security problems that leave it vulnerable to a variety of well-known cyber attacks, any one of which could be catastrophic." Any system for voting over the Internet with common personal computers, they noted, would suffer from the same risks.

The trojans, viruses and other attacks that complicate modern life and allow such crimes as online snooping and identity theft could enable hackers to disrupt or even alter the course of elections, the report concluded. Such attacks "could have a devastating effect on public confidence in elections," the report's authors wrote, and so "the best course to take is not to field the SERVE system at all."

A spokesman for the Department of Defense said the critique overstated the importance of the security risks in online voting. "The Department of Defense stands by the SERVE program," the spokesman, Glenn Flood, said. "We feel it's right on, at this point, and we're going to use it."

An official of Accenture, the technology services company that is the main contractor on the project, said the researchers drew unwarranted conclusions about future plans for the voting project. "We are doing a small, controlled experiment," said Meg McLaughlin, president of Accenture eDemocracy Services.

The Federal Voting Assistance Program, part of the Department of Defense, plans to officially introduce the program in the next few weeks. Seven states have signed up so far to participate: Arkansas, Florida, Hawaii, North Carolina, South Carolina, Utah and Washington. As many as 100,000 people are expected to use the system this year, and the total eligible population would be about one million.

A move to that larger population of voters is far from certain, Ms. McLaughlin said, and the final system could be very different from the one being used this year. "It will be up to Congress and the states to determine if this gets expanded, and how," she said.

"Without doing these experiments, we won't learn more and we won't learn how to help these folks vote in the future," she said.

Trying to vote overseas can be a frustrating ordeal. And Internet voting makes intuitive sense to Americans who have grown accustomed to buying books, banking and even finding mates online.

But the authors of the report adamantly state that what works for electronic commerce doesn't work for electronic democracy: "E-commerce grade security is not good enough for elections," they wrote. The dual requirements of authentication and anonymity make voting very different from most online purchases, they wrote, and failures and fraud are covered by Internet merchants and credit card companies. "How do we recover if an election is compromised?" they wrote.

The report states, "We recognize that no security system is perfect, and it would be irresponsible and naïve to demand perfection; but we must not allow unacceptable risks of election fraud to taint our national elections."

They said any new system "should be as secure as current absentee voting systems and should not introduce any new or expanded vulnerabilities into the election beyond those already present."

One of the authors of the report, David Wagner, an assistant professor in the Computer Science Division at the University of California at Berkeley, said, "The bottom line is we feel the solution can't be a system that introduces greater risks just to gain convenience."

Although some of the possible attacks may sound far-fetched or arcane, the security experts said that each of them had already been seen in some form out on the Internet.

"We're not making up any theoretical concepts," said Aviel D. Rubin, an author of the report and the technical director of the Information Security Institute at Johns Hopkins University. "These are all things that occur in the wild that we see all the time."

Computers on the Internet have become ever more vulnerable to malicious software that takes over the machines' functions to monitor the users' activities, scan them for private information or press them into service to launch attacks on other computers, to send spam or advertise Internet pornography sites online. "And we're going to use these as voting booths?" Mr. Rubin asked. "It

just doesn't make any sense."

A major American election would be an irresistible target for hackers, and the ability of computers to automate tasks means that many attacks could be carried out on a large scale, the report said.

The authors said the Federal Voting Assistance Program, which runs SERVE, and Accenture, the main contractor, should not be faulted for their work, which they found innovative and conscientious. Secure Internet voting, the panel concluded, is an "essentially impossible task."

In fact, the panel said, "there really is no good way to build such a voting system without a radical change in overall architecture of the Internet and the PC, or some unforeseen security breakthrough. The SERVE project is thus too far ahead of its time, and should wait until there is a much improved security infrastructure to build upon."

The risks inherent in SERVE are likely to cripple any system for Internet-based voting, said Barbara Simons, a technology consultant and coauthor of the report. "It's not just a SERVE thing," she said.

Such concerns are not new. They have formed the basis of several recent studies of Internet voting. A report in 2001 by the Internet Policy Institute, financed by the National Science Foundation, concluded that "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."

David Jefferson, an author of the new report and a computer scientist at Lawrence Livermore National Laboratory in Northern California, also worked on a 2000 report for the California secretary of state that reached similar conclusions. "Nothing fundamental has changed," he said, since that report was written.

"Nothing we've seen makes us think that this can be made secure," Mr. Jefferson said.

In attempting to play down the critique of the system, Mr. Flood of the Defense Department called it a "minority report," since it involved only 4 of the 10 outside experts asked to review the system. But Mr. Rubin, the report co-author, noted that the four authors were the only members of the group who attended both of the three-day briefings about the system.

There is no majority report, since the other six experts have not taken a public stance on the project.

Ms. McLaughlin of Accenture said that the company had contacted the other six members of the outside advisory group and that five of the six said they would not recommend shutting down the program.

One of the other outside reviewers, Ted Selker, a professor at the Massachusetts Institute of Technology, disagreed with the report, saying it reflected the professional paranoia of security

researchers. "That's their job," he said.

Mr. Selker, an expert in the ways people use technology, said security is a less pressing concern than mistakes in registration databases, poor ballot design and inadequate polling place procedures. "Every single election machine I've seen — including the lever machine, including punch card machines, including paper ballots — has vulnerabilities," he said.

A security expert and critic of technologically advanced voting systems who had seen an early draft of the study applauded the group's work. "What I saw convinced me that no one should ever vote on that system," said David Dill, a professor of computer science at Stanford University who has become active in voting technology issues. "I understand the problems that people overseas have voting, especially if they are in the military, and I believe we have to make it a lot easier for them," he said. "But SERVE is the wrong solution."



Issue Brief: Internet Voting and Uniformed and Overseas Citizens Absentee Voters

U.S. Public Policy Council of the Association for Computing Machinery

Executive Summary

The reforms introduced by the Military and Overseas Voter Empowerment (MOVE) Act embody the appropriate use of technology. While the MOVE Act is not a panacea, it offers solutions to many of the pressing challenges facing Uniformed and Overseas Citizens Absentee Voters (UOCAVA) voters, and its operational provisions are within reach at a reasonable resource investment. We also support the paper based kiosk model, which we discuss below, for remote voting.

The MOVE Act

Prior to MOVE, an overseas civilian or uniformed service person (UOCAVA voter) was dependent on the postal service to request a registration form, return the registration form, receive a blank ballot, and return the voted ballot. Mail could take a long time to travel between countries. Further delays resulted if a uniformed service voter's location changed during the process, which is common for military members. Consequently, voted ballots of UOCAVA voters often arrived too late to be counted.

Much of the motivation for the MOVE Act, which was signed into law in October 2009, was to address the problems of military voters. Reforms introduced by MOVE include:

- allowing UOCAVA voters to request and receive voter registration and absentee ballot applications electronically;
- requiring states to make blank ballots available electronically at least 45 days prior to any Federal election;
- requiring states to make Federal Write-In Absentee Ballots available online;
- providing for free expedited mail service for voted ballots of overseas uniformed service voters;
- forbidding notarization requirements; and,
- providing UOCAVA voters with the ability to track their ballots.

These reforms dramatically reduce the time required for the entire voting process, while increasing the time available for voters to cast their ballots.

Although MOVE allows for the establishment of one or more pilot programs “to test the feasibility of new election technology,” pilot programs are not mandated by MOVE. In particular, MOVE does not require any pilot program that allows internet or fax voting.

What is Internet Voting?

By internet voting, we mean returning an electronic form of a voted ballot over the internet using email, a web application, or an internet-based fax or phone (e.g. the iPhone). Internet voting can be done from a personal PC, iPhone, cybercafe, library, or kiosk containing a computer dedicated to voting. Kiosk voting can be unsupervised or supervised by authorized personnel; it can be paperless, or it can generate a paper ballot or record that is sent to a local election official. We refer to kiosk voting as dedicated voting; all other forms of internet voting are undedicated.

A major challenge of internet voting (and of any form of electronic voting) is that there is no known way to confidently audit electronic voted ballots, including ballots generated by email, fax, or phone voting. This is because of a fundamental difference between voting and commerce. While fraudulent transactions occur in commerce, we eventually detect them, because commercial transactions create records that are checked by the people who are allegedly the originators of the transactions. Election theft is much harder to detect, because there is only one transaction per person, and that person has no way to later audit his or her vote.

If no reliable post-election audit or recount is conducted, then incorrect software or malicious code could result in the wrong candidate being declared the winner.

One way to provide both the flexibility of electronic voted ballot delivery and the security and confidence that audits provide is to deliver the voted ballot electronically when it is cast and to capture the ballot in a paper form. This process is used by optical scan voting system where a voter marks a paper ballot and scans it into a precinct count optical scanner.

Unfortunately, it is not possible to duplicate this process for most internet voting. Home-based internet voting could create a paper ballot that could be mailed in, but a post-election audit or recount of an internet-based election would require that all internet voters mail in their paper ballots – an almost impossible condition.

Security Risks of Undedicated Internet Voting

Cybersecurity threats make all forms of undedicated internet voting insecure and vulnerable to election theft.

According to a December 2009 report from the Computer Security Institute, a survey of 443 companies and government agencies found that 64% had reported malware infections (malicious software such as viruses or worms) in the preceding year.

The Zeus virus, which steals money from on-line financial accounts, is an especially pernicious example. Because Zeus can simulate the victim's financial statement, the victim will learn of the theft only when some financial transaction cannot be finalized because of insufficient funds. A Zeus-like virus could be created that would steal a person's vote, instead of his money.

Zeus is hardly unique. The Conficker worm has the capability of "calling home" for more instructions. In other words, a Conficker-like virus or worm could remain quiescent until around Election Day, at which time it could call home to find out how its master wants it to vote. If a voter's computer is infected with malware that does nothing most of the time, there is a good chance that the voter will not know that the computer is infected. And even if he is suspicious, detecting and removing malware can be challenging, especially for a non-expert.

Election stealing software on a voter's computer can cast a ballot independent of the voter's intention, and the voter will never know. The computer screen will accurately reflect the voter's choice, but the malware can modify the voter's vote before it is sent over the internet. In other words, it is the malware that votes, not the voter.

Even the companies that produce internet voting software or process internet-based votes are not safe. News reports of government and corporate sites being hacked are becoming more frequent. In a March 2010 talk, FBI Director Robert Mueller is quoted as saying that the FBI's computer network had been penetrated and that the attackers had "corrupted data." The same article discussed the recent successful Google attack, which targeted Google intellectual property, as well as Gmail accounts of Chinese human rights activists:

"Researchers investigating the Google attack -- thought to have affected at least 100 companies including Intel, Adobe and Symantec -- say that prime targets of the hackers were the source code management systems used by software developers to build code."

The implication of Mueller's comments and the Google attack is that voting system software could be rigged by outsiders, including attackers from another country. (The Google attack appears to have originated in China). Another disturbing aspect of the attack targets is that Symantec, one of the targeted companies, is a major supplier of anti-virus and anti-spyware software. The attacked companies, which employ large numbers of computer security experts, have vastly more resources than the relatively small internet voting vendors.

While external risks from hackers is significant, insider risks should not be ignored. As demonstrated by rogue trader Jerome Kerviel, charged with losing about \$7 billion in unauthorized transactions at Société Générale by exploiting his insider status, a malicious trusted insider can be a major threat. Such an insider could hide election-stealing software in large software programs used by vendors and web-based voting sites. If the malware were cleverly hidden, detection could be very difficult.

Other risks of undedicated internet voting include spoofing (creating a fake voting website that looks very much like the real one), phishing (phony email that looks legitimate and attempts to trick the voter into going to an election-rigging website), denial of service attacks, coercion, and vote buying/selling.

The kiosk model has a dedicated computer connected via the internet (presumably using a secure network connection) to a central computer. The computer sends the voted ballot over the network.

Risks of Paperless Kiosk Voting

Paperless kiosk voting has many of the same risks as undedicated internet voting. These include malware infections, denial of service attacks, coercion, and threats to the voter's privacy that jeopardize his right to a secret ballot. While the risks of phishing attacks and vote selling are significantly reduced by the use of a kiosk, the accuracy and security threats, including the possibility that an insider might rig the machine, make paperless kiosk voting unacceptable.

Paper-Based Kiosk Voting.

There are a few forms of paper-based kiosk voting. In one, the voter makes his choices on a dedicated computer that prints out a voted ballot that the voter can verify. If the voter determines that the paper version is not accurate, then both the paper and the electronic versions are voided, and the voter votes again. If the paper version is accurate, then the electronic ballot is cast and the paper ballot is deposited in a ballot box and, together with the other ballots, transported back to the U.S. and ultimately to the appropriate local election official. Even if the computer deployed in the kiosk has incorrect or election-stealing software, the voter can check that the paper ballot or record correctly represents his vote.

Simply generating paper records provides no assurance that it matches the electronic record to which it should correspond; a statistically meaningful audit should be performed to verify the electronic record. If a significant discrepancy is noted, choices are to use the paper record as definitive, use the electronic record as definitive, or



discard the results. Based on our prior analyses, the paper record will provide a more accurate result in most cases.

Alternatively, the kiosk printer can print the appropriate blank ballot, which is then hand-marked by the voter. As above, the ballot will be deposited in a ballot box and transported back to the U.S., where it will be available for audit or recount.

Regardless of how the paper ballots are produced, any pilot program must address chain of custody issues to guarantee that the ballots are securely and promptly transported back to the U.S. in time to be properly tabulated.

There is a risk of vote-buying or coercion in an unsupervised kiosk or a kiosk with supervisors from only one party. Traditional election procedures manage this risk by ensuring that polling places have at least two supervisors with adversarial relationships (such as members of different political parties). Human supervisors are more trustworthy than an unattended computer for authenticating the voter.

Conclusion

We support both the aggressive pursuit of MOVE provisions for internet delivery of election materials to voters and the rapid return of voted ballots, so that they will be counted and available for post-election audits and possible recounts. We also recommend that if internet voting pilots are to be deployed, they should be limited to supervised, dedicated, paper-based systems, coupled with a statistically meaningful audit that should be performed to verify the electronic record.

While returning voted ballots over the internet could improve access and responsiveness for UOCAVA voters, internet voting introduces dangerous risks that can allow elections to be undetectably altered by malicious attacks or buggy software. Without paper ballots, it is impossible to conduct a post-election audit or recount of the internet votes.

Elections are a fundamental component of our national security, and they must be treated as such. Introducing new voting methodologies into real elections demands rigorous risk assessment to ensure the most fundamental election property: integrity.

Association for Computing Machinery (ACM)

With over 90,000 members worldwide, the Association for Computing Machinery is the world's largest educational and scientific computing society, uniting computing



educators, researchers and professionals to inspire dialogue, share resources and address the field's challenges. ACM strengthens the computing profession's collective voice through strong leadership, promotion of the highest standards, and recognition of technical excellence. ACM supports the professional growth of its members by providing opportunities for life-long learning, career development, and professional networking.

About the ACM U.S. Public Policy Council

The ACM U.S. Public Policy Council (USACM) serves as the focal point for ACM's involvement with U.S. government organizations, the computing community and the U.S. public in all matters of U.S. public policy related to information technology. Supported by ACM's Washington, D.C., Office of Public Policy, USACM responds to requests for information and technical expertise from U.S. government agencies and departments, seeks to influence relevant U.S. government policies on behalf of the computing community and the public, and provides information to ACM on relevant U.S. government activities. USACM also identifies potentially significant technical and public policy issues and brings them to the attention of ACM and the community. USACM publishes a monthly newsletter, the ACM Washington Update, which reports on activities in Washington that may be of interest to those in the computing and information policy communities, and highlights USACM's involvement in many of these issues. USACM is actively engaged in number of public policy issues of critical importance to the computing community.

For more information about USACM, please contact the ACM Office of Public Policy at (202) 659-9711 or see <http://www.acm.org/usacm/>.

A comment on the May 2007 DoD report on Voting Technologies for UOCAVA Citizens

David Jefferson¹, Avi Rubin², Barbara Simons³

We have reviewed the Department of Defense report titled "Expanding the Use of Electronic Voting Technology for UOCAVA Citizens" of May 2007. We find the report quite troubling.

Although the report describes many laudable ways to simplify voting for overseas Americans, it also appears fundamentally to be advocating for "a complete Internet voting system", i.e. one that allows voters to cast their ballots on their own PCs and transmit them to the home jurisdiction over the Internet. The report estimates that it would take between 24 and 60 months to develop such a system, depending on recommendations and guidelines.

In 2003 the Department of Defense engaged our services to review its SERVE Internet voting project. The project was subsequently killed because of the numerous and fundamental security problems with it that we documented in a report we issued in 2004 (<http://www.servesecurityreport.org>). We are concerned that this new report appears to be trying to persuade readers that SERVE was a successful project and that Internet voting can be made safe and secure. Unfortunately, it does not accurately reflect the degree of concern that we and many others have expressed about Internet voting.

The new report includes (page 12) *only* the following selective quote from our report:

We want to make it 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 skills 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.

These are about the only lines in our entire report that were not critical of the SERVE project. Those comments were intended to soften an otherwise harsh assessment, and to make it clear that it was the technology, rather than the people, that we were criticizing. The immediately following sentences from our report were not quoted, but they more accurately reflect the report as a whole:

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, and the goal of a secure, all-electronic remote voting system, the FVAP has taken on an essentially impossible task. There really is no good way to build such a voting system without a radical change in overall architecture of the Internet and the PC, or some unforeseen security breakthrough.

¹ Lawrence Livermore National Laboratory, d_jefferson@yahoo.com

² Professor of Computer Science, Johns Hopkins University, rubin@jhu.edu

³ IBM Research (retired), Former President, Association for Computing Machinery, simons@acm.org

In fact, no such security breakthrough has occurred, and we remain convinced that there is no way to secure Internet voting. Perhaps that is why the new DoD report resorts in some places to buzzwords instead of substance. For example, the report claims that roaming digital certificates will be used to combat certain threats. While that may sound good to general audiences, the use of such certificates does not address any of the serious problems identified in our SERVE report.

The IVAS system, deployed in 2006, was a modest successor to SERVE. Although it was reviewed favorably in the DoD report, it actually is more insecure than SERVE. IVAS involved email and fax and did not provide any encryption or authentication of ballots. Several parties, including an independent contractor, were in a position to tamper with or destroy ballots before they were received by local election officials. The DoD report cites surveys of local election officials saying that they would use IVAS again. But while such surveys may indicate interest by officials, they say absolutely nothing about whether such a system is actually secure. We believe it is not.

The current Internet and PC architectures are both such highly insecure platforms that it is essentially impossible to develop a secure system for voting in federal elections on them. From time to time some person or company claims to have “solved” the security problems of Internet-based elections. Such solutions typically deal only with some of the easier issues (voter authentication, secure ballot transmission) by using various encryption mechanisms. Invariably, the most difficult vulnerabilities are ignored, defined away, or addressed with ineffective gestures. Such vulnerabilities include insider attacks of various kinds, phishing attacks, DNS attacks, spoofing attacks, viral and backdoor attacks, distributed denial of service attacks, and automated vote buying and selling schemes. The purported mitigations listed on page 12 of the DoD report are examples of ineffective gestures; reading that list makes one wonder if the authors fully understand the gravity and complexity of the security issues.

Most of the security problems with Internet voting are generic to any PC and Internet application, and *fundamentally have no effective solutions*. This is why the majority of all email transmitted over the Internet is spam, and an estimated 50% of all Internet-connected PCs in the world are infected with malicious software, despite more than a decade of effort and immense investment by the world’s high technology companies in trying to fix these problems. It is not just that no solution to the problems of Internet voting has yet been deployed. The real problem is that *no fundamental solution is possible* using the current Internet protocols and the current PC hardware and software platforms. We do not anticipate that the changes in the design of Internet and in PC hardware and software needed to support secure elections will be forthcoming within the foreseeable future, and certainly not within the five year time span contemplated in this report.

In our 2004 report we made the case against the SERVE Internet voting system. However, those arguments actually apply to *any* Internet voting system, and so we repeat them here (in slightly updated form):

- a) Paperless electronic voting systems have been widely criticized elsewhere for various deficiencies and security vulnerabilities: that their software is totally closed and proprietary; that the software undergoes insufficient scrutiny during certification; that they are especially vulnerable to various forms of insider (programmer) attacks; and that they have no voter-verified audit trails (paper or otherwise) that could largely circumvent these problems and improve voter confidence. All of these criticisms apply directly to Internet voting systems as well.
- b) In addition, Internet voting systems have numerous other fundamental security problems

that generally leave them vulnerable to a variety of well-known cyber attacks (insider attacks, denial of service attacks, spoofing, automated vote buying, viral attacks, etc.), any one of which could be catastrophic.

- c) Such attacks could occur on a very large-scale, and could be launched by anyone in the world, from a disaffected lone individual to a well-financed enemy agency outside the reach of U.S. law. These attacks could result in widespread, selective voter disenfranchisement, and/or privacy violation, and/or vote buying and selling, and/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, some of the attacks could succeed and yet go completely undetected. Even if detected and neutralized, such attacks could have a devastating effect on public confidence in elections.
- d) It is impossible to estimate the probability of a successful cyber-attack (or multiple successful attacks) on any one election. But 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 ever, whether the attacker's motive is overtly political or simply self-aggrandizement.
- e) The vulnerabilities we describe cannot be fixed by better design of Internet voting software. They are fundamental in the architecture of the Internet and of PCs and their software. They cannot be eliminated for the foreseeable future. It is quite likely that they will never be eliminated without a wholesale redesign and replacement of much of the hardware and software security systems that are part of, or connected to, today's Internet.
- f) An Internet voting system might appear to work flawlessly in 2008, or whenever it is first deployed, with no successful attacks detected. Unfortunately, but inevitably, a seemingly successful Internet voting experiment in a U.S. presidential election would be viewed by many as strong evidence that Internet voting can be reliable, robust, and secure. Such reasoning is as fallacious as a claim that our cities are safe from "dirty bomb" attacks because we have been living in cities for a long time and no such attack has ever occurred. Any apparently successful election using Internet voting would encourage expansion of the idea in future elections, as well as the marketing of Internet voting systems to jurisdictions throughout the United States and in other countries.
- g) Just because no successful attack is detected does not mean that none has occurred. Unlike military attacks, many cyber attacks, especially if cleverly hidden, would be extremely difficult or impossible to detect, even in cases when they change the outcome of a major election. Furthermore, the lack of a successful attack in one election 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 Internet voting would make the prize of a successful attack more valuable. In other words, a "successful" trial of Internet voting is the top of a slippery slope toward even more vulnerable systems in the future.
- h) We certainly believe that there should be better support for voting for our military and for citizens living overseas. Unfortunately, we are forced to conclude that it would be a very serious mistake to deploy an Internet voting system. Because the danger of successful, large-scale attacks is so great, we reluctantly recommend against any Internet voting until

both the Internet and the world's home computer infrastructure have been fundamentally redesigned.

Compounding these problems, companies selling Internet voting systems almost invariably claim that the software is proprietary, and refuse to permit examination and evaluation of their systems by independent experts. We fully expect that if this project goes forward, whatever company wins the contract will make exaggerated security claims, as others have in the past, and decline to permit independent experts to attempt to verify those claims and publish the results.

We understand the importance of providing military and overseas U.S. citizens with the best possible access to absentee voting. Many of these people are putting their lives on the line to protect our country, and we support many of the measures in the new DoD report that will make voting easier for them. But, we would do them no favor by providing them with a flagrantly insecure and inauditable method of voting. We believe it would be irresponsible to put our democracy at risk by allowing votes to be transmitted over the wide-open and insecure Internet.

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

January 21, 2004

Dr. David Jefferson, d_jefferson@yahoo.com

Dr. Aviel D. Rubin, rubin@jhu.edu

Dr. Barbara Simons, simons@acm.org

Dr. David Wagner, 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 program's web site is <http://www.serveusa.gov/>. While the system is called an *experiment*, it is going to be used to count real votes in the upcoming general elections. 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 50 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 personnel and their 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) DRE (direct recording electronic) voting systems have been widely criticized elsewhere for various deficiencies and security vulnerabilities: that their software is totally closed and proprietary; that the software undergoes insufficient scrutiny during qualification and certification; that they are especially vulnerable to various forms of insider (programmer) attacks; and that DREs have no voter-verified audit trails (paper or otherwise) that could largely circumvent these problems and improve voter confidence. All of these criticisms, which we endorse, apply directly to SERVE as well.
- b) But in addition, because SERVE is an Internet- and PC-based system, it has numerous other fundamental security problems that leave it vulnerable to a variety of well-known cyber attacks (insider attacks, denial of service attacks, spoofing, automated vote buying, viral attacks on voter PCs, etc.), any one of which could be catastrophic.
- c) 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, and/or privacy violation, and/or vote buying and selling, and/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, some of the attacks could succeed and yet go completely undetected. Even if detected and neutralized, such attacks could have a devastating effect on public confidence in elections.
- d) 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 the attacker's motive is overtly political or simply self-aggrandizement.
- e) 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 all be eliminated for the foreseeable future without

some unforeseen radical breakthrough. It is quite possible that they will not be eliminated without a wholesale redesign and replacement of much of the hardware and software security systems that are part of, or connected to, today's Internet.

- f) We have examined numerous variations on SERVE in an attempt to recommend an alternative Internet-based voting 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. We do suggest a kiosk architecture as a starting point for designing an alternative voting system with similar aims to SERVE, but which does *not* rely on the Internet or on unsecured PC software (Appendix C).
- g) The SERVE system might appear to work flawlessly in 2004, with no successful attacks detected. It is as unfortunate as it is inevitable that 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. (The existence of SERVE has already been cited as justification for Internet voting in the Michigan Democratic caucuses.)

- h) Like the proponents of SERVE, we believe that there should be better support for voting for our military overseas. Still, we regret that we are forced to conclude that the best course 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 immediately and not attempting anything like it in the future until both the Internet and the world's home computer infrastructure have been fundamentally redesigned, or some other unforeseen security breakthroughs appear.

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 describe here, and we have been impressed with the engineering sophistication and skill they have devoted to attempts to ameliorate or eliminate them. 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, and the goal of a secure, all-electronic remote voting system, the FVAP has taken on an essentially impossible task. There really is no good way to build such a voting system without a radical change in overall architecture of the Internet and the PC, or some unforeseen security breakthrough. The SERVE project is thus too far ahead of its time, and should not be reconsidered until there is a much improved security infrastructure to build upon.

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). FVAP's mission is to reduce voting barriers for all citizens covered by the Uniformed and Overseas Citizens Absentee Voting Act (UOCAVA). UOCAVA covers Uniformed Services: U.S. citizens who are members of the Uniformed Services and their family members, and Overseas Citizens: U.S. citizens who reside outside the United States. Uniformed Services are defined as the U.S. Armed Forces (Army, Navy, Marines, Air Force and Coast Guard), merchant marine, commissioned corps of the Public Health Service and the National Oceanic and Atmospheric Administration.

A group of experts in computerized election security, called the Security Peer Review Group (SPRG) was assembled by FVAP to help evaluate SERVE. The task was to identify potential vulnerabilities the system might have under various kinds of cyber-attack, to evaluate the risks these attacks 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 of the SPRG 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. The authors of this report consist of the subset of the SPRG that attended both meetings. Many issues and design improvements were proposed and accepted in those meetings; this report concentrates 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 whether 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 is designed to allow UOCAVA voters both to register to vote in their home districts, and also to vote, entirely electronically via the Internet, from anywhere in the world. Although it is referred to by FVAP as an *experiment*, and is thought of that way by its developers, it is important to realize that it is not just a trial or mock voting system. SERVE will be a complete, medium-scale, federally-qualified and state-certified voting system front end, and it will collect real votes.

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 (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 stored on the central web server for later download by the LEO, at which point 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 stores it for later canvass.

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 voter's PC, 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 voter's machine.

Besides being restricted to overseas voters and military personnel, in 2004 SERVE will be limited to people who vote in one of 50 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.) However, one goal of SERVE is to determine if a similar system might be suitable for expansion in the future to all overseas voters. The total number of eligible overseas citizens plus military and dependents is estimated at about 6 million 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. For these reasons we analyze SERVE not as an experiment, but as a real voting system whose use 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 as applying to SERVE were mentioned in the VOI report. 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.)

VOI also ignored 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". In other words, VOI made no attempt to defend against some of the most serious attacks to which it was vulnerable.

However, the VOI report expressed concern about security problems with "remote" Internet voting (i.e., voting from any Internet-connected computer, anywhere in the world, as permitted under SERVE), and it explicitly declined 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, denial of service, and other 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 (whose assets were acquired in 2003 by

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."*¹

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 Why security for Internet voting is far more difficult than for e-Commerce

Many people mistakenly assume that since they can safely conduct commercial transactions over the Internet, that they also can safely vote over the Internet. First, they usually underestimate the hazards of

¹ The NWI report added a footnote stating "However, remote Internet voting may be appropriate in the near-term for special populations, such as the military and government employees and their dependents based overseas. Such exceptions should be evaluated on a case-by-case basis." However, the viruses, worms, Trojan horses and spyware that we've seen in recent years are much worse than those prevalent at the time of that report, and many of the concerns over insider fraud and the need for voter-verified audit trails have only been understood and articulated since then.

online financial transactions, and are unaware of many of the risks they take even if they are careful to deal only with “secure” web sites through the SSL protocol. But they also assume that voting is comparable somehow to an online financial transaction, whereas in fact security for Internet voting is far more difficult than security for e-commerce. There are three reasons for this: the high stakes, the inability to recover from failures, and important structural differences between the requirements for elections and e-commerce.

First, high security is essential to elections. Democracy relies on broad confidence in the integrity of our elections, so the stakes are enormous. We simply 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 public elections.*

Second, securing Internet voting is structurally different from—and fundamentally more challenging than—securing e-commerce. For instance, it is not a security failure if your spouse uses your credit card with your consent; it is routine to delegate the authority to make financial transactions. But it is a security failure if your spouse can vote on your behalf, even with your consent; the right to vote is not transferable, and must not be delegated, sold, traded or given away. Another distinction between voting and e-commerce is that while a denial of service attack on e-commerce transactions may mean that business is lost or postponed, it does not de-legitimize the other transactions that were unaffected. However, in an election, a denial of service attack can result in irreversible voter disenfranchisement and, depending on the severity of the attack, the legitimacy of the entire election might be compromised.

Third, the special anonymity requirements of public elections make it hard to detect, let alone recover from, security failures of an Internet voting system, while in e-commerce detection and recovery is much easier because e-commerce is not anonymous. In a commercial setting, people can detect most errors and fraud by cross-checking bills, statements, and receipts; and when a problem is detected, it is possible to recover (at least partially) through refunds, insurance, tax deductions, or legal action. In contrast, voting systems must not provide receipts, because they would violate anonymity and would enable vote buying and vote coercion or intimidation. Yet, even though a voting system cannot issue receipts indicating how people voted, it is still vital for the system to be transparent enough that each voter has confidence that his or her individual vote is properly captured and counted, and more generally, that everyone else’s is also. There are no such requirements for e-commerce systems. In general, designing an Internet voting system that can detect and correct any kind of vote fraud, without issuing voters receipts for how they voted, and without risking vote privacy by associating voters with their votes, is a deep and complex security problem that has no analog in the e-commerce world. For these reasons, the existence of technology to provide adequate security for Internet commerce does not imply that Internet voting can be made safe.

1.4 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?” We recognize that no security system is perfect, and it would be irresponsible and naive to demand perfection. However, since we must not allow unacceptable risks of election fraud to taint our national elections, we must have some set of criteria for deciding what risks are acceptable.

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, or sold, and that the voter not be coerced or have his/her privacy compromised under any credible threat scenario, even if the attacker has significant 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 absentee voting for UOCAVA citizens. 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 has eliminated its precinct polling places, so all votes there are absentee. Since the goal of SERVE is to facilitate absentee voting, arguably the security level of absentee voting should be the baseline against

which SERVE is 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 current absentee voting systems.*

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 or coercing voters, 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 offices. In both cases, there are numerous procedural and legal controls that, if adhered to, place perpetrators 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) attack to compromise many votes. It must be essentially *impossible* that any such large-scale attacks 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 such an 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 is, we believe, 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 report.

1.5 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 particular classes of voters could be disenfranchised, based on the likelihood of their voting a particular way. 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. 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 invalidating the voter's ballot by adding too many additional votes. As we shall see, however, electronic voting creates new opportunities for automated, widespread ballot modification that had not previously existed.

A third threat is the loss of privacy—the undermining of the secret ballot. Voting at a properly-managed precinct polling station on paper ballots that are mixed with others in a physical ballot box is the best protection for privacy. The privacy of a voting station or booth is protected by not allowing two people to enter a polling booth together, even if they want to. (An exception may be made for people with disabilities who are unable to vote unaided). Since the ballot is immediately mixed with other ballots in the box, it is virtually impossible to reconstruct who cast which ballot within a precinct. It is much harder to protect the absentee voter's privacy when voting is done using paper absentee ballots filled out at home or at work and sent through the mail, and harder still when an electronic absentee ballot is processed by a large amount of software on several different computers.

A fourth threat, a well-known type of voter fraud, is that a voter may vote more than once. Quoting from the Florida Department of Law Enforcement: Those inclined to do so can capitalize on others' access to an absentee ballot by voting their ballot for them, often with the actual voter not knowing what has occurred. This offers tremendous opportunity for vote fraud, particularly to those who have access to the ill or infirm or those who do not have the ability to resist the influence of another as they are urged to vote in a

“required” manner. It also encourages those inclined to commit voter fraud to seek to utilize absentee ballots provided to those whose interest in voting is marginal or non-existent.²

As in the case of ballot modification, Internet voting provides new opportunities for multiple voting by allowing those willing to invest time in social engineering to determine those registered voters who are most unlikely to participate in an election. After obtaining that information, identity theft would allow Internet voting to be automated so as to achieve a significant amount of false participation in any given election.

A fifth threat, intrinsic to absentee voting, is vote buying, selling, and trading. 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 is also easy to use computers 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.

1.6 Vulnerabilities in SERVE

Lack of voter-verified audit trail and insider attacks. DRE (direct recording electronic) voting systems have been widely criticized because they are essentially unauditible. First, there is no way that a voter can verify that the vote recorded inside the machine is the same as the vote that he or she entered and saw displayed on the machine’s touchscreen. And later, if serious problems occur in the canvass of the votes (which happens all too frequently with DREs), there is no independent audit trail of the votes to help resolve the problem.

These issues have been debated widely around the nation in the last couple of years, and the computer security community, including all of the authors of this report, nearly unanimously supports the position that no DRE voting system should be permitted that does not have some sort of voter-verified audit trail. Voter verification is the only readily available effective defense against programmed insider attacks. Since so much has been written about this subject we will not repeat the arguments here. For further information see, for example, Prof. David Dill’s web site www.verifiedvoting.org, or Prof. Rebecca Mercuri’s site www.notablessoftware.com/evote.html, or the California Secretary of State’s report and directives on touchscreen voting at www.ss.ca.gov/elections/touchscreen.htm. What we wish to note here is that every argument about the need for voter verification and auditability that have been made about DRE systems also apply essentially unchanged to the SERVE system. Indeed, from a voter’s point of view, SERVE can be said to act as one giant DRE machine.

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 and signature on the outer envelope are checked against the voter list. Assuming they match those of a registered voter who has not yet voted, the outer envelope is opened in the presence of at least two people. 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

² http://www.fdle.state.fl.us/publications/voter_fraud.asp

is available to everyone. In each voting district that is participating in SERVE, a local election official (called the 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 before transmission, if someone were to intercept the encrypted ballot en route, it would be impossible to decipher the actual vote. 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. SERVE decrypts the ballot using the SERVE private key, separates the ballot from the voter's name, and then encrypts the ballot (without the voter's name) using the LEO's public key. Therefore, only the appropriate LEO will be able to decrypt the encrypted ballot. The encrypted ballot is stored 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 could deduce how voters in his/her precinct have voted by downloading votes from SERVE so frequently that they get at most one new vote and voter name each time. Recall that 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 to infer 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 and with access to the SERVE database 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 the 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 were to land in the wrong hands and old system keys were exposed. Mathematical breakthroughs in the future, for example, could expose old keys.

In Sections 2 and 3, we describe in detail several further risks to voter privacy. While today's absentee voting systems have their own privacy risks, taking all these risks as a whole, we believe that SERVE heightens the risk of large-scale privacy compromise.

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 by allowing vote buyers to automate the process. 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, what's new about SERVE is that 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. 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. Finally, it is possible that many servicemembers would share the same computer, so obviously, those votes would legitimately come from the same IP address.

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. In short, the possibilities for large-scale, automated vote buying, selling, and swapping in SERVE go beyond anything present in existing absentee systems.

Intimidation. Intimidation is a potential problem with all forms of absentee voting, since the voter is not guaranteed the privacy of the voting booth. But it can be even more of a problem with Internet voting if the voter is not using his/her personal machine, since the owner of the machine may have installed software that would record what the voter is doing.

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, terrorists, 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.

Undetectable 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, 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. The mere presence of a confirmation screen does not prove

that the vote was recorded correctly. These concerns are analogous to those that have been expressed about DRE's without a voter-verified audit trail.

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 re-holding 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. Many states have 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 derive revenue from 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 pop-up 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.

1.7 Organization of This Report

The remainder of this report is organized as follows. The technical core of the report analyzes three significant threats to the security of SERVE: attacks made possible by lack of control over the voting environment (Section 2), web spoofing and man-in-the-middle attacks (Section 3), and denial of service attacks (Section 4). Finally, we present our conclusions and recommendations (Section 5). Also, several appendices discuss additional issues, including other security risks of SERVE (Appendix A), a prior experience with Internet voting (Appendix B), a possible alternative to SERVE (Appendix C), and fundamental problems with writing secure and bug-free software (Appendix D).

The security threats to SERVE are summarized in Table 1, where we characterize 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	There are a million ways to make a complex transaction such as voting 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	medium	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 to go 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 attacker	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.³ 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, 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 plug-ins, screen savers, and Microsoft 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.

Voters' home computers 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.

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 contains a malicious payload designed to take control of machines and wreak havoc with a future election. Since virus checking software programs defend against only previously known viruses, virus checkers often are unable to keep up with the spread of new viruses and worms. Consequently, malicious worms are widespread among Internet-connected computers today. For instance, in 2001, the Code Red worm infected 360,000 computers in 14 hours, and in 2003 the Slammer worm brought down many ATM machines and compromised many Internet hosts.⁴ Modern worms are even more virulent, are often spread by multiple methods, are able to bypass firewalls and other defenses, and can be difficult to analyze. For example, it took quite a while to realize that SoBig.F was a large-scale Trojan horse designed to plant spam engines.⁵

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." New viruses almost certainly 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

³ http://www.amanet.org/research/pdfs/ems_short2001.pdf

⁴ <http://securityresponse.symantec.com/avcenter/venc/data/w32.sqlexp.worm.html>

⁵ <http://securityresponse.symantec.com/avcenter/venc/data/w32.sobig.f@mm.html>

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 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 damage 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 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*.⁶ 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 one party's partisan to prevent a large fraction of the other party from voting, while leaving most or their own party unaffected. 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,

⁶ 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.

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 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 attack could be implemented, 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. The attacker would be able to see all of the traffic by decrypting and re-encrypting 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 re-sign the patched ActiveX control; becoming an ActiveX signer requires fooling one of the certifying authorities, or simply purchasing a key. Browsers come with over one hundred default keys that are already completely trusted, whether the users know it or not.

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 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 web server 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.⁷ It was later discovered that these damaging attacks had been perpetrated by a lone teenager not on US soil.⁸

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).⁹ In 2003, an Internet election in Canada was disrupted by a denial of service attack on Election Day.¹⁰ 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 web server 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 web server's computational

⁷ <http://www.nipc.gov/investigations/mafiaboy.htm>

⁸ <http://www.cnn.com/2000/TECH/computing/04/18/hacker.arrest.01/>

⁹ <http://www.symantec.com/avcenter/venc/data/codered.worm.html>

¹⁰ http://cbc.ca/stories/2003/01/25/ndp_delay030125