

How Security Problems Can Compromise Remote Internet Voting Systems

Guido Schryen

Institute of Business Information Systems
RWTH Aachen University
Templergraben 64
52062, Aachen, GERMANY
schryen@winfor.rwth-aachen.de

Abstract: Remote Internet voting systems still suffer from many security problems which rely on the clients, the servers, and the network connections. Denial-of-service attacks and viruses still belong to the most challenging security issues. Projects and studies like the “Voting Technology Project” of CALTECH and MIT or SERVE of the US Department of Defense set up to gain experience evidence many of the notional weaknesses of current Internet voting systems.

1 Introduction

Theoretical research about the security of electronic voting systems started many years ago and countless approaches have been proposed since then. Not only motivated by academical research, but also quickened up by the US-presidential election’s dilemma in 2000 several practical projects were conducted to assess the feasibility of electronic voting systems over the Internet. But reducing election problems to the counting process itself – as it might happen due to the big election in 2000 – clouds some more issues to be faced. How many votes have been destroyed, how many eligible voters have been disenfranchised from voting, how many votes have been altered in the context of absentee voting? Most people trust in the established offline voting procedures and show little interest in security issues as long as computers and networks are not involved. Actually, the real extent of election fraud is undetected, only some are known and published. The report of CALTECH and MIT [CM01, p.3] mentions: “*Our data show that between 4 and 6 million votes were lost in the 2000 election.*” Jefferson et al. [Je04, p.11] report: “*A recent example [of election fraud] involved boxes of paper ballots that were found floating in San Francisco Bay in November, 2001.*”

These incidents alone strongly motivate the discussion of the use of Internet voting systems and their ability to successfully address election fraud. Furthermore, supporters of these systems argue that there will be a higher voter turnout and more trust in elections. But unfortunately, using the Internet with its current architecture and protocols would cause more security trouble than we can handle.

The paper is about this trouble and the Internet's inappropriateness for remote voting scenarios. Section 2 shows the differences to e-commerce systems and discusses security aspects concerning the voting clients, voting servers, and the network connections between them from a theoretical point of view. Supplementary, section 3 summarizes Internet voting reports of some of the most important projects and links these experiences to the insights gained in sec. 2. Finally, conclusions are drawn in sec. 4.

2 Security problems

Security issues of Internet voting systems can be discussed from many points of views, e.g. technology driven, political science driven, or judicial driven. I address this field with a technology view, focussing especially on voting servers, voting clients, and the network infrastructure enabling the client-server-connections.

2.1 Differences to e-Commerce

Sometimes it is assumed by mistake that safely conducting commercial transactions over the Internet with SSL and server-side certificates means that one can also safely vote online using the same mechanisms. However, this is wrong, as Internet voting is different in many aspects [Je04]:

- Elections are inseparably linked to democracy and malfunctioning election processes can directly and decisively influence it. Democracy relies on broad confidence in the integrity of elections. Consequently, Internet voting requires a higher security level than e-Commerce does.
- It is not a security failure if your spouse uses your credit card with your consent, but the right to vote is usually¹ not transferable.
- A denial-of-service (DoS) attack might occur and prevent you and others from performing e-Commerce transactions. But generally there is a broad time window and after detecting and fixing the DoS attack business can be transacted. In the context of Internet elections a DoS attack can result in irreversible voter disenfranchisement and the legitimacy of the entire election might be compromised. For example, voters who want to cast their ballot during the last minutes of the voting time window would have no other voting channel available.
- Business transactions require your authentication by sending passwords, PINs, or biometric data. Voting however, requires authentication only when you register for an election and when you cast your ballot due to authorization, but concurrently demands anonymity to the vote (decision). This implies the adoption of much more complex security protocols.

¹ Exceptions must be allowed for blind and other handicapped people.

People can detect errors in their e-Commerce transactions as they have audit trails: they can check bills and receipts and when a problem appears recovery is possible through refunds, insurance, or legal action. Vote receipts (showing the vote decision and proving that the vote was unalteredly counted) must not be made out, as otherwise votes can be paid and extortion might occur.

2.2 Assumptions and focus

I consider only those voting scenarios whose voting protocols base on public-key-cryptography, certificates, and a public key infrastructure without addressing the protocols itself detailed, but this is no strong constraint. Furthermore I assume the potential voters to use ordinary PCs with Windows or Linux software and an arbitrary connection to the Internet.

Technological security issues are to be found in several dimensions (see figure 1, for a more detailed discussion see [Sch04]), but below I focus on hardware, software, and infrastructure as some of the most critical issues from my point of view. Voting protocols aren't less important but are basically out of range of this article.

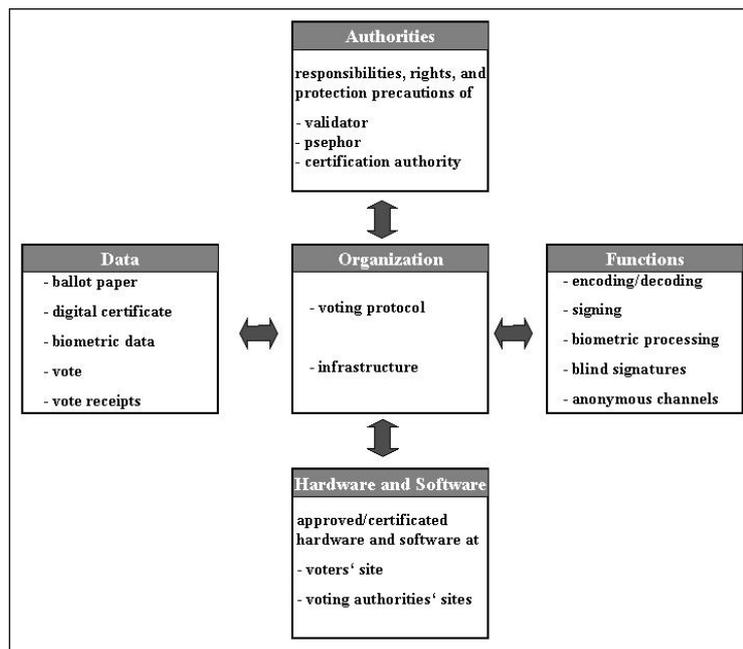


Figure 2: Security dimensions for voting systems [Sch04, p.7]

The following subsections address security issues of the client, the (voting) servers, and the connections between clients and servers. In particular I look at the voting process itself as opposed to online voter registration, which is a separate, but important and difficult problem.

2.3 Client related security issues

One of the most significant problems clients are facing is malicious payload (programs and configurations). Rubin [Rub02] analyzes this problem: There is virtually no limit to the damage viruses, Trojan horses, sniffing programs, etc. can cause. Although the presence of security defense software (virus and intrusion detection) becomes more and more widespread the current state of the art does often not go much beyond comparing a program against a list of signatures. If the security software vendor hasn't updated his definition files due to unknown signatures e.g., then a computer might remain unprotected for a while including the voting window. The option that the malicious payload and its signature will not be detected makes it all even worse. Using trusted software in the sense of signing software by a trustworthy entity and checking the digital signature of programs sounds like a sustainable concept, but this means that each piece of software has to be signed and checked. First, there is no software or hardware architecture supporting this, and secondly, Jefferson et al. [Je04] report cases where people were tricking Microsoft into signing a malicious ActiveX control. Summing up today there is no foolproof test for whether or not malicious payload is installed.

Rubin [Rub02] mentions the software Back Orifice 2000 (BO2K) that is freely available and fully open source tool for remote control of a computer. Once it is installed on a machine, it enables a remote administrator (or attacker) to view and control everything on that machine. As it is open source, an attacker might change the code so that it remains undetected by security defense software (due to a new signature). As it runs in stealth mode even a sophisticated administrator would have difficulties to detect it. Voting decision could be read, changed, and blocked from being sent without discovery.

As election dates are known in advance the activation of malicious software can be effectively triggered. The Chernobyl virus for example was scheduled for April 26, 1999, and affected many computers by modifying the BIOS in such a way that they couldn't even boot. If that happens on the day of an election many eligible voters would be disenfranchised. Politically ambitious attackers could target a particular demographic group aiming at a direct effect on the election's result.

And even worse it does not take a very sophisticated malicious payload to disrupt an election, as easy web browser attacks demonstrate. Most common browsers come with an option for a proxy setting that indicate that all web communications should take place via a proxy; the proxy is interposed between the (web) client and the (web) server and completely controls all Web traffic between these two. The proxy option can be easily changed by just adding a few lines to the preference file. Using the Netscape browser you just change the file `prefs.js` by adding these lines indicating that all web traffic goes to the corresponding server and port:

```
user_pref("network.proxy.http", www.malory.com);  
user_pref("network.proxy.http_port", 1799);
```

Although proxies cannot be used to read information in a secure connection, they can be used to spoof a user into a secure connection with the attacker, instead of the actual voting server.

Unfortunately, there are many ways for attackers to attach malicious payload to common PCs, most of us have probably experienced at least one option.

- Malicious payload can be installed by having physical access to the computer. Administrators in companies have full privileges on many computers and can infect them using setup routines on floppy disks and CDs. Many more scenarios are possible granting full physical access to an attacker.
- Most common malicious code is distributed via emails. Think about Melissa, I Love You, Sobig.F, and MyDoom/Novarg which infected probably millions of computers in a very short time. You don't even have to open an email attachment to get infected, e.g. the virus Bubbleboy was triggered as soon as a message was previewed in the Microsoft Outlook mailer. We can observe an alarmingly increasing activity.
- Buffer overflows are a known and well used point of attack. This kind of attack occurs when a process assigns more data to a memory location than was expected by the programmer. Web server programs and web browsers have proved to be susceptible for buffer overflows when arbitrary attacker's code can be executed. Buffer overflows are one of the most common form of security flaws in deployed systems today.
- A widely accepted but also dangerous way of executing programs is the use of ActiveX controls which are native code residing on the web server and attached to web content. If your browser's settings allow ActiveX controls to be executed they are automatically and maybe unknowingly downloaded and started. Trojan horses can be installed that way and on day of election brought to attacking execution. Many people use ActiveX controls as browser plug-ins, screen savers, calendars, etc., consciously or not. ActiveX controls can perform as man in the middle. This attack together with spoofing is addressed in the next subsection.
- Vendors of widely spread software like graphic programs, word processing program, etc. are in a strong position to change software and configuration files while the setup process is running. On day of election the changes can compromise or bother the voting process on this machine. Just let one rogue programmer of the software vendor be interested in subverting an election.

Authentication in the context of a public key infrastructure is done by signing data with the private key. Assumed the voter has a private key it must not be stored on the hard disk, floppy disk, CD, or USB stick, but should be kept on a secure key store like a smart card. As smart card readers are not directly connected to voting servers (voting) data flow through the insecure PC environment where it can be changed or blocked. Blocking of votes is easy: malicious code ensures that the vote gets not forwarded to the voting server.

Changing the vote is possible when you actually sign other data than you intended to sign: While your computer's display makes you believe you sign your vote for party A the malicious code changes your vote in favor of party B and sends this to the card reader. If this reader has no dedicated display allowing to double-check the vote then the voter might be fooled. The attacker doesn't even have to know your private key. Consequently, card readers without a(n) (expensive) display are insecure in this sense. Most voting systems don't even integrate any kind of card readers as they are not widely spread.

Today, mobile devices as voting clients drop out [IPI01, p.16]. Beside technical security problems displays are still limited in terms of display area, color, and resolution, as well as text input capability. They may easily be lost or stolen, and the cost for providing these devices to registered voters could be prohibitive.

Rubin [Rub02] sums it up: "In current public elections, the polling site undergoes careful scrutiny. Any change to the process is audited carefully, and on election day, representatives from all of the major parties are present to make sure that the integrity of the process is maintained. This is in sharp contrast to holding an election that allows people to cast their votes from a computer full of insecure software that is under the direct control of several dozen software and hardware vendors and run by users who download programs from the Internet, over a network that is known to be vulnerable to total shutdown at any moment."

2.4 Server related security issues

The problem of DDOS attacks affects all participating servers. In this section we focus on the voting servers but generally the considerations can be applied to all servers. Attacks where legitimate users are prevented from using a system by malicious activity, are known as denial-of-service-attacks (DOS attacks). If many attacking machines collaborate to mount a joint attack on the target machine we talk about a distributed DOS attack (DDOS attack). In this scenario, an attacker could take control of many computers (called "zombies" or "slaves") in advance by spreading a virus or worm, and the slaves are waiting for instructions of a master computer to blindly follow them. There are mainly two forms of (D)DOS attacks: (1) The adversaries swamp the network connection of the targeted server with junk data that clogs up the network and prevents other, legitimate traffic from getting through. The SYN flood attack that exploits a weakness of the Internet protocol TCP is a famous example. (2) The adversaries are able to overload the server's computational resources with useless tasks that keep it busy. SSL-protected websites are susceptible to this kind of (D)DOS attack as the SSL protocol requires the recipient to perform a slow cryptographic operation (typically an RSA private-key computation).

Suffering a DDOS attack voting servers are in danger of being cut off from the Internet and eligible voters resulting in their disenfranchisement. If DDOS attacks are targeted demographically (regional voting server is attacked) and we have a close voting campaign then they could sway the election. DDOS attacks are huge and real problems and no effective protection mechanism is known.

Many DDOS attacks have occurred, an example of an DDOS attack on domain name servers is reported in the following subsection.

Another (easier) way to target a machine and to make it crashing is the *ping of death attack* [Rub02].

If voting clients would act as DRE (direct recording electronic) voting systems they wouldn't suffer from (D)DOS attacks as they could store the vote and send it later. Unfortunately, this approach seems currently not feasible, because it is not practical or desirable for PCs to emulate all the characteristics of DRE systems² [IPI01].

2.5 Connection related security issues

The sore spot of connection related attacks is the fixed election time window. Attackers can focus the last hours of the election window and paralyze the network of a region that is assumed to vote for candidate A by the majority. Even a quick fixing can take some hours resulting in the disenfranchisement of voters and affecting the election's result. One form of attack affects the Internet's Domain Name Service (DNS). The DNS is used to maintain a mapping from IP addresses, which computers use to reference each other (e.g. 134.130.176.7) to domain names, which people use to reference computers (e.g. www.winfor.rwth-aachen.de). The DNS is known to be vulnerable to attacks. Currently, there are just 13 DNS root server, some big companies additionally mirror them. In 2002 the DNS servers were exposed to a distributed denial-of-service-attack (DDOS) where several servers were fully loaded.³ If on election day the DNS servers aren't available for many voters, then a connection to the vote server is not possible. Only those voters who know the IP address of their voting server could vote then.

Another attack is DNS spoofing where the true IP address of a domain name is overwritten with a fake IP address. The control of DNS root servers might be difficult, but the heavy use of DNS caching (on local or regional servers due to speeding up) makes this impossible. Although answering this problem with the protocol DNSSEC (RFC 2535 und 2931) would be effective, its practical impact is low. Facing DNS spoofing the voter follows the instruction for voting and enters the denoted domain name. But unknowingly he gets a wrong IP address and he is spoofed into a communication with an attacker. He might receive a page that looks like the voting page.

Then the attacker acts as man in the middle giving him the power to abolish votes. The same happens in the context of social engineering: an attacker sends emails to voters containing links to the attacker's computer. When they look authentic many people would trust this email. Theoretically, this kind of spoofing can be effectively addressed with digital certificates of web sites, but today most people are not familiar at all with SSL connections and certificates and hence wouldn't check or discover this fraud.

² For more information about DRE systems visit <http://www.verifiedvoting.org/drefaq.asp>.

³ <http://www.washingtonpost.com/ac2/wp-dyn?pagename=article&node=&contentId=A828-2002Oct22¬Found=true>

Similar attacks could also work against the registration process. Eligible voters could be let 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 they were not registered. This could exclude them from voting.

Not to forget are attacks on Internet router which forward IP packets through the Internet to the server and back. If IP routers fail due to DDOS attack a whole region might be unable to cast votes.

Some attacks could be mitigated with the existence of a vote receipt proving that your vote arrived. As this receipt must not contain the vote decision⁴ (see discussion above) itself it just proves that a vote decision arrived. There is no guarantee of data integrity, i.e. your vote could have been changed on your computer, on a computer in the network, or on the voting server. Many DRE (direct recording electronic) voting systems don't have any sort of voter-verified audit trail. Furthermore, how can you be sure that your vote was actually counted and not left behind? Traditional elections don't feature this problem as the whole process can be peered (except for absentee balloting).

3 Internet Voting Reports

Some projects have been set up to scrutinize the appropriateness of the Internet for a remote voting system. The most important ones are the *Voting Technology Project* of CALTECH and MIT [CM01], *A Report on the Feasibility of Internet Voting* of the California Internet Voting Task Force [CV00], *A Security Analysis of the Secure Electronic Registration and Voting Experiment (SERVE)* [Je04], the *National Workshop on Internet Voting* of the Internet Policy Institute [IPI01], and *i-vote* of the Research Group Internet Voting [IV02].

Most projects come (after a detailed security discussion) to the conclusion that today the Internet should not be used for remote voting as the architecture, protocols, hardware, and software feature many vulnerabilities that could easily allow attackers to compromise elections. Only the German study [IV02] looks a bit more optimistical on Internet elections. Two projects [CV00; IPI01] distinguish between several stages of Internet voting and concede practicability for supervised Internet voting clients. The following subsections summarize the results of the corresponding reports.

⁴ The Internet Policy Institute [4, p.19] discusses an approach that provides voters with the ability to vote multiple times, and have only the last vote count. However, some practical problems arise and make this concept difficult to be implemented.

3.1 CALTECH and MIT: Voting Technology Project

The CALTECH/MIT Voting Technology Project was initiated academically and conducted by the California Institute of Technology and the Massachusetts Institute of Technology as an interdisciplinary approach. It is not restricted to Internet voting scenarios.

However, regarding Internet voting they find [CM01, p.15; 42]: *“However, Internet voting, in the judgment of many experts, is not ready for wide-scale use. There are three problems. First, there are concerns of coercion if Internet voting is done from remote locations, such as the voter’s home computer. Second, large-scale fraud is more likely because it is easier to hack the entire system if it is on the Internet, than it is to coordinate many millions of voters voting at precincts or thousands of poll workers. Third, many people do not have computers at home or are sufficiently intimidated by computers that Internet voting (either from home or at the precinct) might create a further obstacle to voting for millions of voters. [...] Delay Internet voting until suitable criteria for security are put in place.”*

3.2 California Internet Voting Task Force: A Report on the Feasibility of Internet Voting

The California Internet Voting Task Force was convened by Secretary of State Bill Jones to study the feasibility of using the Internet to conduct elections in California.

They define four steps of Internet voting and propose an evolutionary approach where stages 1 and 2 feature a supervised use of an Internet voting machine and stage 3 and 4 integrate remote Internet voting: (1) Internet Voting at Voter’s Polling Place, (2) Internet Voting at Any Polling Place, (3) Remote Internet Voting From County Computers or Kiosks, and (4) Remote Internet Voting from Any Internet Connection.

The opinion of the Task Force is [CV00, p.1f]: *“At this time, it would not be legally, practically or fiscally feasible to develop a comprehensive remote Internet voting system that would completely replace the current paper process used for voter registration, voting, and the collection of initiative, referendum and recall petition signatures. [...] However, current technology would allow for the implementation of new voting systems that would allow voters to cast a ballot over the Internet from a computer at any one of a number of county-controlled polling places in a county. [...] The success or failure of Internet voting in the near-term may well depend on the ability of computer programmer and election officials to design a system where the burden of the additional duties placed on voters does not outweigh the benefits derived from the increased flexibility provided by the Internet voting system.”*

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

The SERVE voting system was built for the U.S. Department of Defense's FVAP (Federal Voting Assistance Program) [DoD01] and intended to be deployed in 2004 for U.S. citizens living overseas; participating states are Arkansas, Florida, Hawaii, North Carolina, South Carolina, Utah, and Washington. In the meantime the Pentagon refused to deploy the system in 2004 due to strong security concerns [DoD04]. A heavy security discussion was triggered by the security analysis report conducted by independent scientists. They disclosed that the SERVE voting system suffers from most security risks discussed above, stating [Je04, p. 3]: *"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."*

Surprisingly, without any security discussion it was announced that overseas voters can still vote by fax [DoD04].

3.4 Internet Policy Institute: National Workshop on Internet Voting: Issues and Research Agenda

The National Workshop on Internet Voting was funded by the National Science Foundation (NSF) and conducted by the Internet Policy Institute and the University of Maryland. It was former President Clinton who requested the NSF to examine the feasibility of online (Internet) voting.

Internet voting systems are grouped into poll site systems where voting machines are placed in traditional polling places, kiosk systems with voting machines located in convenient locations as malls, libraries, and schools, and remote systems where any computer that is Internet accessible might serve as a voting machine.

The core conclusion is [IPI01, p. 23]: *"Poll site Internet voting appears potentially able to meet currently accepted levels of risk; remote voting, however, does not, at least with current or soon available technology. The possibility of large-scale automated attacks on remote Internet voting systems leads to a level of risk so high as to be unacceptable."*

3.5 Research Group Internet Voting : i-vote

The German Research Group Internet Voting of the University Osnabrueck has conducted a project including the set-up of an Internet voting system and evaluating it empirically in the context of real elections. The report doesn't criticize remote Internet elections in principle, but argues more fuzzily claiming absolute secure voting clients, the certification of voting software and voting systems, and the use of chip cards with digital signatures. It admits, too, that much security research still has to be done.

4 Conclusions

Remote Internet voting heavily struggles with security issues and possible attacks that arise from the infrastructure, protocols, hardware, and software. There remain not only conceptual questions like how to deal with voting receipts and which voting protocol to use, but also everyday Internet problems like Trojan horses, viruses, spoofing, DDOS attacks, etc. Most reports clearly decline the appropriateness of today's Internet for remote elections. Two characteristics impose security stakes on a level we haven't faced before: (1) Remote Internet elections technically open a former closed voting environment to attackers all over the world who can gang together to selectively strike election processes. (2) The impact of a disrupted election can be large: the whole election might be questioned by an unsettled society and not less worse the election result might be notelessly effected. As our societies and states base on democracy and sound elections no described security risk is tolerable. According to Rivest [Riv01] adopting remote electronic voting means that we would have sacrificed too much security for the sake of voter convenience. However, the scale of security measures depends on the meaning of the election: voting a student parliament is not comparable with voting a national parliament that rules a state. Furthermore, supervised voting terminals and a closed Internet voting infrastructure don't feature many problems discussed above and are worth being more explored.

References

- [CM01] California Institute of Technology (CALTECH) and Massachusetts Institute of Technology (MIT): Voting Technology Project, 2001. Available at <http://www.vote.caltech.edu/>
- [CV00] California Internet Voting Task Force: A Report on the Feasibility of Internet Voting, 2000. Available at <http://www.ss.ca.gov/executive/ivote>.
- [DoD01] US Department of Defense: Federal Voting Assistance Program, 2001. Available at <http://www.fvap.gov/index.html>.
- [DoD04] US Department of Defense: Pentagon Decides Against Internet Voting This Year. American Forces Information Services News Article, Feb. 6, 2004. Available at http://www.defenselink.mil/news/Feb2004/n02062004_200402063.html.
- [Je04] Jefferson, D.; Rubin, A.D.; Simons, B.; Wagner, D.: A Security Analysis of the Secure Electronic Registration and Voting Experiment (SERVE), 2004. Available at <http://www.servesecurityreport.org>.
- [IPI01] Internet Policy Institute: Report of the National Workshop on Internet Voting: Issues and Research Agenda, 2001.
- [IV02] Research Group Internet Voting: i-voteReport: Chancen, Möglichkeiten und Gefahren der Internetwahl. Zusammenfassung der Ergebnisse und Empfehlungen der Forschungsgruppe Internetwahlen zur Nutzung des Internets für Wahlen, 2002.
- [Riv01] Rivest, R.: Electronic Voting, 2001. Available at <http://theory.lcs.mit.edu/~rivest/Rivest-ElectronicVoting.pdf>.
- [Rub02] Rubin, A.: Security Considerations for Remote Electronic Voting over the Internet. Communications of the ACM 12 (45), pp. 39-44, 2002.
- [Sch04] Schryen, G.: Security Aspects of Internet Voting. Proceedings of the 37th Hawaii International Conference on System Sciences, 2004. Available at <http://csdl.computer.org/comp/proceedings/hicss/2004/2056/05/205650116b.pdf>.