Security Considerations for Remote Electronic UOCAVA Voting

Nelson Hastings
Rene Peralta
Stefan Popoveniuc
Andrew Regenscheid

NIST

National Institute of Standards and Technology

U.S. Department of Commerce

[This page intentionally left blank.]

NISTIR 7770

Security Considerations for Remote Electronic UOCAVA Voting

Nelson Hastings
Rene Peralta
Stefan Popoveniuc
Andrew Regenscheid
Information Technology Laboratory
National Institute of Standards and Technology
Gaithersburg, MD 20899-8930

February 2011

[This page intentionally left blank.]

TABLE OF CONTENTS

This page intentionally left blank.]

1 Introduction

The Election Assistance Commission (EAC) requested that the National Institute of Standards and Technology (NIST) research technologies to improve uniformed and overseas United States citizens' ability to vote, as required by the Uniformed and Overseas Citizens Absentee Voting Act (UOCAVA) [1]. Additionally, the Help America Vote Act of 2002 (HAVA) requires the Technical Guidelines Development Committee, with technical support from NIST, to study remote access voting, including voting over the Internet [2]. This report contains the results of NIST's research into threats and security technologies related to remote electronic voting for overseas and military voters.

In December 2008, NIST released NISTIR 7551, *A Threat Analysis on UOCAVA Voting Systems* [3], which documents the threats to UOCAVA voting systems using electronic technologies for all aspects of overseas and military voting. NISTIR 7551 considered the use of postal mail, telephone, fax, electronic mail, and web servers to facilitate transmission of voter registration materials, blank ballots, and cast ballots. It documented threats and potential high-level mitigating security controls associated with each of these methods. The report concluded that threats to the electronic transmission of voter registration materials and blank ballots can be mitigated with the use of procedures and widely deployed security technologies. However, the threats associated with electronic transmission, notably Internet-based transmission, of cast ballots are more serious and challenging to overcome and the report suggested that emerging trends and developments in that area should continue to be studied and monitored.

While NISTIR 7551 looked at a variety of technologies for all aspects of the UOCAVA voting process, this report takes a deeper look specifically at the issues associated with remote electronic voting over the Internet. It identifies and defines desirable security properties of remote electronic voting systems and major threats faced by these systems that could violate those security properties. It also discusses the current technologies that could be used to mitigate some of those threats and open issues that may still need to be addressed.

In August of 2010, the EAC posted their *UOCAVA Pilot Program Testing Requirements* document [6]. This document defines requirements for remote electronic voting systems using a supervised-kiosk architecture that is intended for use in a UOCAVA pilot program. However, this report considers all remote electronic voting systems, with particular attention to the threats and technologies for remote voting from personally owned and operated devices. Depending on how it is used, the supervised kiosk model mitigates

many of the threats identified in this document, particularly those related to software integrity, coercion, vote-selling, and voter identification and authentication.

1.1 Purpose and Scope

On April 26, 2010, the EAC submitted their *Report to Congress on EAC's efforts to Establish Guidelines for Remote Electronic Absentee Voting Systems* [7], detailing a roadmap intended to be used by the EAC, NIST, and the Federal Voting Assistance Program (FVAP) to create and implement guidelines for remote electronic absentee voting systems for overseas and military voters. The initial phase of this roadmap calls for a report describing security issues related to remote electronic absentee voting system for UOCAVA voters. This report, along with NIST's initial report on threats to UOCAVA voting systems, NISTIR 7551, *A Threat Analysis on UOCAVA Voting Systems* [3], is intended to meet this need.

This document is part of a series of documents that address the UOCAVA voting. In addition to NISTIR 7551, NIST has released drafts of NISTIR 7682, *Information Systems Security Best Practices for UOCAVA-Supporting Systems* [4] and NISTIR 7711 *Security Best Practices for the Electronic Transmission of UOCAVA Election Materials* [5]. In addition to NIST's research on security issues associated with remote electronic UOCAVA voting, NIST is also researching usability and accessibility topics. A report documenting this research, *Accessibility and Usability Considerations for Remote Electronic UOCAVA Voting*, will be released in early 2011.

1.2 Intended Audience

This document is intended for election officials, technologists, advocacy groups, UOCAVA voting system vendors, and other members of the elections community that will be working with the EAC, NIST, and the FVAP on improving the UOCAVA voting process with the use of electronic technologies. While this document assumes familiarity of the UOCAVA voting process and a high-level understanding of information system security technologies, it is intended to be accessible to a wide audience.

1.3 Organization

The remainder of this report is organized as follows:
- **Section 2** provides a high-level description of the remote electronic voting system architectures that are analyzed in the remaining sections this document. The primary architecture considered is remote voting over the Internet from personally-owned devices.
- **Section 3** provides an overview of the structure for the sections containing the subtopics: Confidentiality, Integrity, Availability, and

Identification and Authentication. Each subtopic contains a discussion of the potential benefits, properties, threats, current and emerging technical approaches and open issues.

- **Section 4** discusses issues related to confidentiality of remote electronic voting systems. Confidentiality refers to the concept of ballot secrecy, and also to protecting sensitive voter information and system data from unauthorized disclosure. This section discusses desirable properties of remote voting systems to deal with confidentiality issues, threats, and possible mitigating technologies.
- **Section 5** discusses issues related to integrity of remote voting systems. This includes data integrity, aimed at safeguarding important election records, including cast ballots and audit logs, as well as software integrity. It describes desirable properties of systems intended to support data and software integrity and identifies threats and possible technical approaches for dealing with these issues.
- **Section 6** describes properties, threats and technologies related to availability of voting systems. Availability refers to the ability of the system to be ready for use when needed by voters and election officials in the face of malicious and incidental threats.
- **Section 7** discusses issues related to the identification and authentication of voters, system operators, election officials, and system components. It identifies threats to the authentication process and discusses various technical methods for authenticating users and components.
- **Section 8** summarizes the important findings report.

2 General Architecture

The following section provides an architectural view of remote voting systems in order to provide a reference from which to discuss security considerations presented in the rest of the document.

2.1 System Components

The general architecture of electronic remote voting systems, as shown in Figure1, is composed of several different components. The following subsections detail the components that may be found in an electronic remote voting system.

2.1.1 Voters' Platforms

Figure 1 shows three different platforms that may be used by a voter to request, receive, and cast their ballot: personal computers, public computers, and kiosks.

Figure 1- General Architecture for Remote Electronic Voting
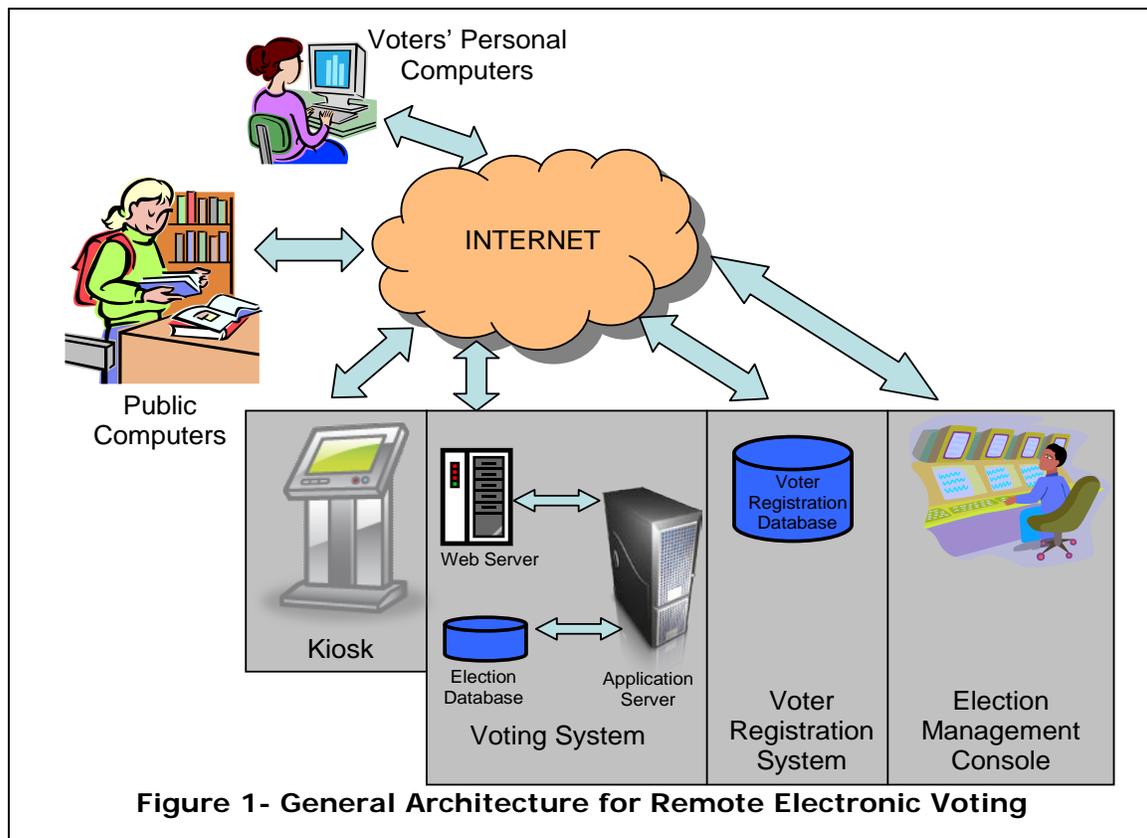
Personal computers refer to general purpose computing systems a voter may have at home for their personal use, including desktop and laptop computers, tablets, and smart phones. Voters may also use general purpose computer systems found at public locations such as libraries, schools, and Internet cafes and are referred to as public computers. Finally, voters may use dedicated devices called kiosks that may or may not be under the control and supervision of poll workers and/or election officials. In general, the voter's platforms will have a connection to the Internet in order to complete the voting process.

The voter's platform is not under the control of election officials except in a supervised kiosk voting system architecture. This means that there may be no poll worker or election official to ensure the voter's privacy has not been compromised or that voters have not been coerced into casting their ballot differently than they desired. In addition, the platforms not under the control of election officials may be poorly protected and vulnerable to malware, phishing, and denial of service attacks. These platforms may be the target of attacks to monitor and/or modify voter choices, capture personal information, or prevent a voter from accessing the voting services.

2.1.2 Voting System

Figure 1 shows the voting system consisting of three subcomponents: web, database, and application servers. This is a simplified representation of the three subcomponents since they may include other hardware and software not shown in the diagram to ensure system reliability and availability.

The web server provides the interface that voters use to interact with the remote electronic voting system. The web server interface may have the voter use a general purpose browser or a voting-specific client application to obtain voting services from the voting system. The web server has a connection to the Internet so voters can interact with the remote electronic voting system. In addition, the web server will interact with the application server that provides the voting services to the voter.

The application server contains the logic for the services provided by the remote electronic voting system. The services provided by the application server may include the ability for the voter to: register to vote, request a blank ballot, return completed ballots, tally the ballots, and generate election reports. The application server has an indirect connection to the Internet via its interactions with the web server. This provides the voter interface to the remote electronic voting system. In addition to interacting with the web server, the application server will interact with election database and possibly the voter registration system.

The election database contains the ballots for the different jurisdictions serviced by the remote electronic voting system. When a voter requests a ballot, the application server queries the election database to find the appropriate ballot for the voter based on their information. In addition, the election database server may store completed ballots when they are not stored on the application server. The election database server usually does not have a direct connection to the Internet. Access to the database takes place through the application server.

In general, the web server, application server, and election database are housed in one location, such as a data center managed by a jurisdiction or commercial third party. The locations that house the servers and database will need to provide the physical storage space, communication connections, and physical and logical security measures.

2.1.3 Voter Registration System

Voter registration systems are run by states and contain a repository of eligible voters who can participate in elections. The voter registration system assembles the repository of eligible voters using information from different sources such as department of motor vehicle records, judicial records, and possibly the remote electronic voting system. States provide jurisdictions with the registered voter information when elections are held. Jurisdictions can use the information to ensure that only eligible voters are allowed to cast ballots and that only one ballot is cast per voter. Figure 1 shows the voter registration system being accessed directly via an Internet connection or a more limited connection such as a state or military operated network. The jurisdictions may use their connection to the voter registration system to access the voter information in real-time during the election or to make electronic copies of the information they need at a given point during the election.

2.1.4 Election Management Console

Election officials that administer elections use the election management console. The election management console provides an interface to the voting system so administrative task, such as the configuration of ballots, defining the time and date to cast ballots, setting up the tallying rules for the election contests, and the generation of election reports, can be completed. The election management console can be located in the same place as the voting system or may be at some other location (such as the office of the election officials).

2.1.5 Component Connectivity

In general, the components that voters interact with (e.g., voters' personal computers, public computers, and kiosks) use the Internet as their connection to the voting system.

Remote electronic voting system servers and other backend system components may be on the same local network or connected to one another over the Internet.

2.2 Authorized Users

Each of the components of a remote electronic voting system is under the control of one or more different people called users. The users that control the different components are authorized to perform certain, but possibly not all, actions on the component. Although the users are authorized to perform actions on the components, they have the potential to attack the remote election voting system. This section will describe the different users found in the remote electronic voting system but will leave the description of the potential threats which these users present for Section 2.3 Threat Sources.

2.2.1 Voters

The basic voting functionality required by a voter is to: (a) submit voter registration information, (b) request and receive blank ballots, (c) complete a ballot, and (d) return a completed ballot. Voters may use their own personal computers, public computers, and/or kiosks to interface with the remote electronic voting system. In general, voters only have limited capabilities on public computers and kiosks.

Kiosks typically do not have general-purpose applications, such as word processors or email clients, so voters do not have access to these types of applications when voting from a kiosk. However, public computers may provide voters with access to applications other than voting, such as word processors, email clients, and web browsers.

When using their own personal computers, it is the responsibility of the voter to install, configure, and protect their personal computers and the applications that reside on the computer. The different platforms voters use to interface with the voting system have different security and function advantages and disadvantages when considering remote voting system architectures.

2.2.2 Election Officials

Election officials require the capability to administer an election, including adding or removing voters from the voter registration database, configuring ballot styles, defining the time and date to cast ballots, setting up the tallying rules for the election contests, and the generation of election reports. Election officials may interface with the remote electronic voting system via the election management console. As described in Section 2.1.4, the election management console may or may not be co-located with the voting system.

2.2.3 System Administrators

System administrators will require the capability to install, configure, and protect the different components of the remote electronic voting system. In addition, the system administrator will ensure the components they are responsible for can connect to other components of the remote voting system as needed. The system administrator will monitor the components they are responsible for to look for signs the components are operating improperly or are under attack. The system administrator will vary from component to component. Depending on how the architecture is implemented, third party service providers may make up the system administrator for all the components except for the voter's personal computers. Voters are the system administrators for their personal computers. Election staff will serve as system administrators for the kiosk, voting system, voter registration system, and election management console.

2.2.4 Auditors / observers

Auditors and observers will need access to information generated or observed during an election in order to perform their functions. In general, auditors and observers will have limited information collected through observation due to the distributed nature of remote electronic voting systems. Most of the information auditors and observers will have access to will be electronically generated by the remote electronic voting system with a possible exception when paper ballots are used or a voter verified paper audit trial is produced. The integrity and accuracy of the information used by the auditors and observers will greatly impact the effectiveness of their functions.

2.3 Threat Sources

Threat sources are groups or individuals that could feasibly attack a voting system. Some attacks on voting systems could be conducted by almost any dedicated individual, while others may require significant resources, knowledge or access to voting system equipment. Threat sources can be

broken down into two classes: internal and external sources. Internal sources are individuals or groups with some level of authorized access to the voting system equipment or the supporting infrastructure (e.g. the communications network). External sources are individuals or groups that do not have any special level of authorized access to the voting system equipment or supporting infrastructure. This report considers the following examples of threat sources.

2.3.1 Internal Threat Sources

In general, internal threats come from individuals or organizations with privileged and authorized access to the remote electronic voting system required to support or carry out use of the system in an election. Threats from inside sources may be more dangerous and more difficult to protect against since they have some level of access to the system.

Voters: Voters' access to the remote electronic voting system is limited through the voters' platform used: their own personal computers, public computers, and kiosks. In general, voters will not have direct access to the voting system, voter registration system, or election management console. Voters are allowed to submit voter register information, request and receive blank ballots, complete a ballot, and return a single completed ballot. However, voters may use their voting platform to try to cast multiple ballots using multiple credentials, prove how they voted to sell their vote, expand their access to damage the voting system, change the results of the election, or harm the credibility of the election results.

In addition, the voting platforms may pose a threat to the remote electronic voting system without the voters' knowledge or cooperation. When voting platforms contain malware, the voting platform may try to inhibit a voter from casting his or her ballot, alter a voter's choices, monitor how a voter votes, use the voter's credential to gain and expand access to damage the voting system, change election results, or harm the credibility of the election results. Although the voter is not actively participating in attacking the remote electronic voting system, the platform they use to interact with the voting system poses a threat that appears to be from the voter.

Election Officials: Election officials access the remote electronic voting system via the election management console and possibly voting system equipment as authorized users on the voting system component. Election officials are allowed to add eligible voters to the voter registration database, remove ineligible voters, configure ballot styles, define the time and date to cast ballots, set up the tallying rules for the election contests, and generate election reports. However, election officials may not need to be able to

install and configure applications or have unrestricted access to the remote electronic voting system equipment. Election officials will have access to election data, such as cast ballots and system event logs, on the remote electronic voting system that most other authorized users may not. Access to the election data may allow a malicious election official to modify the results of the election, monitor how people vote, and provide incorrect ballot configurations.

Similar to the voter and voters' platform, the election official and election management console may pose a threat to the voting system without the election official's knowledge. If the election management console contains malware, the console may try to prevent ballots from being cast, alter ballot configurations, monitor how voters vote, and use the election official's credential to gain and expand access to damage the voting system, change election results, and harm the credibility of the election results. Although the election official is not actively participating in attacking the remote electronic voting system, the console they use to interact with the voting system poses a threat that appears to be from an election official.

System Administrators: System administrators access the remote electronic voting system equipment via a remote connection or a terminal directly connected to the equipment. In addition, system administrators have physical access to the equipment. System administrators are allowed to install, configure, and monitor the remote electronic voting system equipment to ensure the equipment is functioning properly. System administrators may directly administer the components of the remote electronic voting system or the supporting infrastructure used by the system. For example, network technicians at telecommunication companies or Internet Service Providers (ISPs) are system administrators of the infrastructure used by the remote electronic voting system. Election IT staff are system administrators for the election management console when it is located at the election official's office. System administrators have a level of access to the system that no other authorized user has in order to configure and maintain the system. Given this level of access, system administrators may try to prevent ballots from being cast, alter ballot configurations, monitor how voters vote, damage the voting system, change election results, or harm the credibility of the election results.

Other insiders: There are other internal individuals or organizations that may have access to the remote electronic voting system equipment before, during, or after an election cycle. For example, voting system manufacturers will have access to the software source code and hardware designs used to implement their remote electronic voting system. This level of access provides an opportunity for errors to be introduced, maliciously or not, into

the components of the remote electronic voting system. Voting system integrators have similar access as voting system manufacturers, but without access to the software source code or the designs of hardware components. This level of access provides the opportunity for known software and hardware errors to be exploited, and for third party, non-voting specific software and hardware to be integrated into the remote electronic voting system components containing errors; malicious or not. The support staff of different organizations, including but not limited to jurisdictions, voting system manufacturers, voting system integrator, and third party service providers, may have access to the remote electronic voting system equipment and that provides an opportunity for the system to be exploited. Examples of support staff include administrative assistants, package and mail delivery personnel, and warehouse personnel.

2.3.2 External Threat Sources

In general, external threat sources come from individuals or organizations not needed to support or carry out use of the system in an election.

Hostile Individuals: Individuals and affiliated individuals may attempt to inhibit ballots from being cast, monitor how voters vote, damage the voting system, change election results, and harm the credibility of the election results. These individuals rely on their technical knowledge and ability to deceive legitimate users and administrators. In general, attacks from hostile individuals are limited based on resources – time, money, and people – they can accumulate or control as required for a given attack scenario.

Hostile Organizations: Like hostile individuals, hostile organizations that may not have legitimate access to the remote electronic voting system in order to attempt to inhibit ballots from being cast, monitor how voters vote, damage the voting system, change election results, and harm the credibility of the election results. Hostile organizations can marshal more resources, particularly money and people, to conduct an attack on the remote electronic voting system than an individual. Given these added resources, a hostile organization can recruit, hire, and train individuals, as well as obtain more costly technology to conduct an attack on the system. Hostile organizations can take many forms including civilian, foreign-sponsored, or terrorist organizations.

3 Overview

The remainder of this report discusses security issues that need to be considered when developing, deploying, or using remote electronic voting systems. The discussion divides the issues into four topic areas:

- ***Confidentiality:*** Confidentiality refers to the concept of ballot secrecy and also the protection of sensitive voter information and system data from unauthorized disclosure. Issues related to confidentiality are discussed in Section 4.

- ***Integrity:*** This includes data integrity, aimed at preventing important election records, including audit logs and cast votes, from being improperly modified, as well as software integrity. Issues related to voting system integrity are discussed in Section 5.

- ***Availability:*** Availability refers to the ability of the system to be accessible to voters and election officials in the face of malicious and incidental threats. Issues related to voting system availability are discussed in Section 6.

- ***Identification and Authentication:*** Identification and authentication includes the identification and authentication of voters, system operators, election officials, and system components. Issues related to the identification and authentication of voting system users and components are discussed in Section 7.

These areas were chosen to break the discussion of security issues into closely related topic areas. Issues related to any one of these topic areas are closely bound to those associated with other topics. For instance, an insufficient authentication mechanism could allow an unauthorized individual to access sensitive information (a confidentiality violation) or modify key voting system records (an integrity violation).

For each topic area, this report discusses the following:

- ***Potential Benefits:*** The move from the current mail-in absentee voting process to a remote electronic voting system can provide some benefits to security, such as in the areas of automated forms of strong authentication, timeliness of delivery, and ballot secrecy. For each of the topic areas, this report will describe the advantages of remote electronic voting.

- **Properties:** In order to facilitate discussion of threats to remote electronic voting systems, this report provides lists of desirable security properties. In general, threats identified in this report are actions that can violate one or more of those properties. The security properties identified in this report are based on properties and requirements identified in other electronic remote voting system documents including the Secure Electronic Registration and Voting Experiment (SERVE) Project documentation [10], the Common Criteria Protection Profile for online voting systems [8], and the Council of Europe's standards for online voting systems [9]. Policymakers ultimately must decide which properties must be met by voting systems to be acceptable in their jurisdictions. This report provides notes with each property that can help policymakers decide which properties are realistically achievable with current and emerging security technologies.

  This report provides definitions for the identified desirable security properties. While definitions may be written in absolutes, readers should recognize there are always tradeoffs that have to be made. For example, the extent a security property can be met versus the cost and usability of implementing the property. Acceptable tradeoffs must be made when deploying systems which often necessitates compromising strict interpretations of some of the proposed properties.

- **Threats:** This report describes some of the major threats to remote electronic voting systems. However, this document is not intended to be a thorough threat or risk assessment on remote electronic voting systems. This document describes some of the more serious threats to remote electronic voting systems. It does not attempt to enumerate all threats. Readers should consult other resources, such as NISTIR 7551, for information on additional threats.

- **Current and Emerging Technical Approaches:** This report identifies and describes some existing and emerging technologies that can be used to mitigate some of threats faced by remote electronic voting systems.

- **Open Issues:** Some security issues associated with remote electronic voting do not have complete solutions at this time. In some instances, advances in technology are needed to address threats, while in other cases the technology is developed, but is not widely deployed.

4 Confidentiality

Voting systems must protect the confidentiality of sensitive information stored on those systems. Notably, remote electronic voting systems have unique concerns about protecting ballot secrecy compared to polling place systems. While an electronic voting machine in a polling place typically does not learn the identities of voters interacting with it, remote electronic voting systems typically must identify and authenticate voters in order to verify their eligibility and provide them with the proper ballots. In some jurisdictions, local or state election procedures dictate that the identities of overseas and military voters must be able to be linked to cast ballots, a property usually forbidden in polling place systems. Despite this, remote voting systems must protect their information from being used illegitimately.

Remote electronic voting systems must also protect the confidentiality of other sensitive information on those voting systems. Remote electronic voting systems may include an online voter registration database containing sensitive personally identifiable information. They must also protect sensitive system information that could be used to compromise the security of the system, such as secret cryptographic keys or passwords.

4.1 Potential Benefits

Compared to mail-in voting, remote electronic voting systems have the potential to provide much greater technical controls for maintaining ballot secrecy. With mail-in voting, ballot secrecy is protected by procedural means: identities of voters are physically separated from cast ballots prior to viewing the contents of the ballots. Small-scale ballot secrecy violations are still possible if colluding election workers handling mail-in ballots do not follow proper election procedures. Access control mechanisms and cryptographic technologies can provide strong protections against attacks on ballot secrecy. Technical measures can be taken so an arbitrarily large number of trusted officials must collude to violate ballot secrecy.

Furthermore, remote electronic voting systems can also provide some protection against unsophisticated attempts to coerce voters. For instance, systems may allow voters to cast multiple ballots and only count the final ballot issued by the voter. If voters feel pressure to vote a particular way in one instance, they would be able to cast a new ballot at some other time or location free from improper influence. While it is significantly more difficult to block coercion attempts from more sophisticated or determined attackers, this is still a useful benefit offered by remote electronic voting systems.

4.2 Properties

This section discusses high-level properties aimed at assuring confidentiality of the vote and of the voter. Confidentiality is necessary to protect the autonomy and privacy of the voter as well as the secrecy of the vote.

A strong form of enforced confidentiality, called receipt-freeness, is also discussed. This property makes it impossible for the voter to prove to a third party how he or she voted. This property addresses the threats of coercion and buying/selling of votes.

Property: Ballot Secrecy

The voting system protects the secrecy of cast ballots.

Notes:

All voting systems leak some information about voters' choices. Such information can usually be derived from data made public during the election (e.g., partial tallies, lists of voters). The remote electronic voting system should not add to this loss of secrecy in any meaningful way. In particular, a voter should not lose plausible deniability regarding his or her vote. Protecting ballot secrecy does not necessarily mean that it must be impossible to link individuals to cast ballots; state law regarding ballot secrecy differs from jurisdiction to jurisdiction. While the general public should not be able to perform this linkage, election officials acting in accordance with state and local election law and procedures may be required to have the capability to link voters to cast ballot. For these cases, voting systems should implement protections to ensure that ballot secrecy can only be breached when proper procedures are followed. For example, the system could force multiple trusted election officials to jointly interact with the system to violate ballot secrecy, and the system could only provide mechanisms for linking single ballots, not all ballots at once.

Property: Protection of Personal Information

The voting system protects voters' personal information from unauthorized disclosure.

Notes:

The voting system should not needlessly store voters' personal information. Any personal information that is stored should be protected against unauthorized disclosure. Use of encrypted storage is recommended in order to minimize the damage caused if storage media is lost or stolen, and access control mechanisms should be used to limit access to sensitive information.

Property: Receipt-freeness

Voters are not able to provide convincing evidence of their ballot selections to third parties.

Notes:

The threat of vote selling and coercion attacks becomes more serious if voters are able to give attackers evidence of how they voted. This information could be used to reward the voter for voting correctly in a vote-selling attack or as evidence that the voter met the demands of a coercer.

Notably, remote voting systems should not increase the likelihood of large-scale buying and selling of votes compared to current mail-in voting methods. They also should not increase the likelihood of large-scale coercion of voters. Coercion is different from vote buying in that the voter is not a willing participant.

Property: Protecting sensitive system data from improper disclosure or use

All sensitive system information handled by the voting system should only be readable by authorized administrators or election officials.

Notes:

Examples of sensitive system data are: passwords or keys used by the election officials to access, configure, and run the voting system; and timestamps recording when voters authenticated or cast ballots.

Property: Minimal storage

The voting system only stores sensitive information necessary to ensure the correct functioning of the voting system.

Notes:

While there are many safeguards that can be put in place, online systems are at risk for unintended data breaches. Internet-accessible systems should not store sensitive information that is not needed by the system. Notably, voter registration databases may contain sensitive voter information, such as identification numbers, that may not be needed by the voting system. When the voting system operates its own voter list or database, sensitive data fields should not be copied over from the primary voter registration database unless the information will be used by the voting system.

Property: Limited communication

Only necessary communications traffic is passed between entities participating in the voting process.

Notes:

As a general rule, there should be limited communications between voting system components. Passing extraneous information, even information that may look benign, increases the chance that this information could be combined to violate confidentiality goals, such as ballot secrecy.

4.3 Threats to Confidentiality

This section discusses some of the more significant threats to confidentiality that are either unique to remote electronic voting systems or that may be more severe in this context. This is a high-level classification that addresses generic threats for all remote voting systems. It does not address threats to individual voting system implementations.

4.3.1 Central System Data Breaches

A data breach is an unintentional release of secure information to an unauthorized party. In the context of voting systems, data breaches can cause loss of vote secrecy as well as loss of private voter information. The potential damage of private information exposure may be less severe in voting systems than in some other systems, such as financial databases or health databases, since voting systems do not need to store as much sensitive private information.

Storage of unencrypted sensitive information carries increased risk and should be avoided when possible. Connection to the Internet also increases the risk of a data breach. Failure to properly secure encryption keys and passwords can result in granting unauthorized access to malicious (or simply curious) third parties. Poor key management can result in insufficiently vetted personnel (e.g., temporary workers) obtaining decryption keys that they are not supposed to have. This can lead to serious data breaches. Additionally, compromised keys can harm the integrity of stored or in-transit data.

A remote electronic voting system may use an external database (e.g., a vehicle registration database). In this case, the voting system could become a route for exposure of private information contained in the external database. Standard database security practices should prevent sensitive information from being exposed. However, the scenario in which two

database administrators each assumes the other is responsible for preventing data breaches is a concern.

4.3.2 Coercion

Voting systems that allow the voter to vote more than once can make it harder to effectively coerce voters (since voters could vote again at a later time). On the other hand, if the secrecy of the vote is not secured, then coercion can be a more serious problem than in non-electronic voting. The reason is that electronic coercion attacks can scale easier and impact more voters and ballots. In particular, coercion that takes the form of reprisals long after the election has ended could be a serious problem, should the secrecy of the vote be compromised on a broad scale. If the voting system has a capacity to link cast ballots to voters (say, under a court order or a voter challenge), then it may be desirable to implement a mechanism for permanent removal of this capacity. In principle, this would occur via destruction of secret keys after a prescribed amount of time has elapsed. Keys that are meant to be eventually destroyed could be split into electronic components and tamper-evident physical components to help ensure the keys are destroyed. In modern information systems, it is very difficult to fully ensure the destruction of electronic data.

4.3.3 Buying and Selling of Votes

A concern with remote electronic voting is the possibility of a market for voting credentials could emerge. A similar threat exists in the case of mail-in voting, in which the unfilled ballots could be bought and sold. However, the scalability and increased anonymity inherent to remote electronic voting potentially makes this a more serious concern. We do not know how to gauge the likelihood of this threat in the presence of law-enforcement deterrents. We note that, in most cases, this threat requires the willingness of both buyer and seller to commit a crime. This should serve as a significant deterrent to vote selling for most of the voting population. On the other hand, any change in voting technology implies a corresponding change in the cost/benefit equations that determine the extent of illegal practices such as vote selling.

A related concern is vote swapping (i.e., vote pairing). This occurred in the 2000 and 2004 elections in the US. It is conceivable that the deployment of Internet voting could cause a surge in this practice if there is an easy mechanism to exchange credentials to voting systems or verify how individuals voted.

Since long-lived voter credentials may increase the likelihood of these types of threat, it may be advantageous to have voters obtain at least part of their voting credentials in the days or weeks prior to the election.

4.3.4 Malicious Software on Client Systems

An emerging threat to computer systems over the last few years is that of malicious software infecting computers, giving attackers control of these systems. Researchers from the Georgia Tech Information Security Center have estimated that attackers may control 15 percent of online computers in this way [12]. What "control" means here is that the machines have been infected by malware that allows some level of access to them. The level of access is typically enough to steal private information and tap communications. Compromised machines could potentially violate the secrecy of the vote. Votes could be linked to machines or, depending on the voting protocol, even to voter identities. While this is clearly illegal, it is unclear what value this information might be to criminals. Unlike credit card numbers, there is no clear financial gain from knowing how a person voted. This is particularly true if such knowledge cannot be verified by a third party (as anyone can claim to know how someone else voted). Furthermore, this type of information is typically only valuable in bulk (as a reference, a single stolen piece of credit card information sells for between $0.85 to $30 [14]). Bulk voting information has two principal uses: tying demographics to voting and large-scale voting coercion. The former is easily obtainable from statistical analysis. The latter seems to be a low-likelihood threat on two accounts: i) it necessitates verifiable information; and ii) it appears hard to do without getting caught.

If compromised machines are able to steal verifiable voting information, then another threat scenario is plausible: vote buying and selling. Opinions vary regarding the severity of the vote buying and selling threat.

4.4 Current and Emerging Technical Approaches

This section discusses the main tools at our disposal for secure implementation of remote electronic voting systems. Some of the tools are standard IT security mechanisms, whereas others are of special applicability to voting.

4.4.1 Cryptographic Protections

Cryptography can protect any data that is communicated from one system to another as well as stored data. For example, the data which travels through the Internet between the voting system and the voter's computer can be efficiently protected from unauthorized access via protocols like

Secure Socket Layer (SSL) or Transport Layer Security (TLS) [15]. SSL and TLS are widely-deployed encryption mechanisms that are often used to protect communications between a web server and browsers. When used with mutual authentication, these protocols provide end-to-end security.

When used to protect data at-rest, cryptographic keys can be split between several people, requiring an arbitrary number of key holders to come together to decrypt data. Such mechanisms offer protection against insider attacks, as long as a small number of insiders can be trusted to not collude in an attack.

Proper cryptographic key management is very important to achieving protection using cryptographic techniques. Keys must be generated, stored, used, and destroyed in specific ways to ensure there are not ways to bypass the cryptographic protections.

4.4.2 Advanced Cryptographic Voting Techniques

Modern cryptology provides several possible solutions for securely conducting secret-vote online elections. These solutions provide very good properties in idealized scenarios where voters make no mistakes, have complete control of their computers, and communication lines are reliable. The scenarios typically allow for fraudulent voters attempting to sabotage the election and for attackers having unimpeded read access to all communication lines. The result of these idealized protocols is that a tally of the votes of all honest voters is obtained and is publically verifiable without compromising the secrecy of the votes.

Despite there being an abundance of voting protocols with the above properties ([16][17][18][19][20] are just a few), the problem of remote voting using the Internet is far from solved. This is because the Internet is not the idealized scenario assumed by that body of work. Voters make mistakes and their computers may be partially under the control of malware. Communication lines may not be reliable. Also, there have been no formal usability or accessibility studies of current cryptographic voting schemes yet, but researchers anticipate that such studies would identify issues that would need to be addressed. Further research may lead to dramatic improvements, but current cryptographic voting techniques do not solve many of the challenges associated with remote electronic voting.

4.4.3 Access Control Mechanisms

Access control mechanisms can be used, in conjunction with identification and authentication mechanisms, to restrict access to data, applications or

actions to particular users. Different levels of access can be granted to different users; a relatively common set of access levels include read, write, and execute permissions, and modern access control mechanisms often provide more fine grain control over permissions. Access control can be implemented in many different ways. On computer systems, access control mechanisms are most often enforced by operating systems, and, in the case of voting systems, voting applications.

For example, access control mechanisms could provide only a designated election official with the access rights to write, modify or delete ballot definition files, but give a much wider set of users access rights to only read those files.

Access control mechanisms could also implement things such as dual-person control, whereby the system requires two or more users to authenticate to the system before providing access to a particular resource. However, such functionality is often not provided by modern operating systems or applications. When used, dual-person control is often implemented with a combination of technical and procedural means.

Depending on how access control mechanisms are implemented, it may be possible to bypass those protection mechanisms. For example, if access control mechanisms are enforced by an application, users may still be able to access resources through the operating system. If the operating system enforces access control mechanisms, an individual with physical access to the system may be able to access resources by booting from a different operating system. Furthermore, in many modern operating systems, the system administrator, or root user, often has nearly unlimited control over the system. For these reasons, it is important to also use cryptographic protections to restrict access to sensitive data, rather than solely relying on common operating system or application-level access control mechanisms.

4.4.4 Separation of Duties

With a combination of procedural and technical means, operators of remote voting systems can enforce separation of duties to limit the capabilities of any single user or computer system. For instance, important information or tasks could be split between several election officials or system operators, requiring them to collude to conduct an attack. One example of how this could be implemented is that one official could be given a key to a locked room with voting system equipment, while a second official is given a credential for administering the voting system equipment.

4.5 Open Issues

Achieving a very strict notion of ballot secrecy remains a challenging issue in remote electronic voting systems. While polling place voting systems do not store, or even learn, the identities of voters, remote electronic voting systems need to authenticate voters before allowing them to cast ballots. Cryptographic protocols exist to protect the secrecy of ballots even from those with unrestricted access to voting system equipment, but these technologies may not be ready for immediate use with remote electronic voting systems. For technical, procedural, and legal reasons, it is likely that any deployed voting system for UOCAVA voters would still have access to, and probably store, sufficient information to violate ballot secrecy. Depending on policy decisions at state and local levels, this issue may not require a technical solution beyond what is already practical.

Advanced voting-specific cryptographic protocols have highly desirable properties in idealized models, but in practice, systems based on these protocols are often difficult to use and require that cryptographic keys be distributed to voters before an election. These systems also do not protect against many types of attacks, particularly if the computer used to cast votes and the voting environment are not secured.

Current techniques for remote electronic voting do not solve the problems of coercion and vote selling that are inherent to unsupervised voting. Variations on these attacks are possible with mail-in absentee voting, although in that voting method, it is difficult for a single individual to impact many voters. When moving to remote electronic voting, election officials and technologists should consider whether the move makes it easier to scale these attacks. In particular, there appear to be ways that attackers could coerce or buy votes remotely. A simple attack involves selling or transferring the credentials that voters use to log into the remote voting system. This particular issue and threat will be discussed further in the Identification and Authentication section (Section 7).

Despite IT professionals' and users' best efforts, data breaches continue to occur, releasing personally identifiable information (and other sensitive information) to attackers. This problem is not unique to voting systems. For the time being, it may be impossible to guarantee the secrecy of voter information stored on voting system equipment from determined and technically sophisticated attackers. However, there appears to be very little reason to store potentially valuable sensitive information on these systems. Depending on the type of information stored by the voting system, there may be very little motivation to attempt to illegitimately access this information.

5 Integrity

This section discusses security issues associated with voting system integrity. Integrity refers to the trustworthiness of the system, including both the data on the system and the functions provided by the system's software. Maintaining integrity involves implementing safeguards to ensure data and software on a system are not modified by unauthorized parties. It is typically preferable to have these safeguards block unauthorized attempts to modify data or software, but in some cases, it is only possible to detect integrity violations.

Integrity includes the concept of the origin or source from which the integrity is based upon. In other words, the origin or source of the integrity for data or software functionality can be traced back to a particular trusted authoritative entity. Tracing integrity back to a particular entity is closely related to identification and authentication, which is covered in Section 7.

5.1 Potential Benefits

5.1.1 Authenticity of Electronic Records

A cryptographically signed record of each cast ballot can be issued by the voting system components and transmitted for tallying and auditing purposes. The signed record can be easily and exactly replicated to reduce the likelihood of data loss. Assuming adequate key management, the signed record cannot be forged. Authenticity can be verified using public key cryptography.

5.1.2 Strong Integrity Protections In-Transit

It is a common misconception that the greatest threat associated with conducting transactions over the Internet is the modification of information as it is being transmitted. While this is a potential threat that must be mitigated, in fact there are very good technical solutions for protecting information during transmission. Cryptographic protocols, such as TLS or Internet Protocol Security (IPSec), are very effective at providing integrity protection in-transit.

5.2 Properties

There are two main categories of properties for integrity: data integrity and software integrity. Data integrity is related to the integrity of the election records, especially those records directly used to derive the final election tallies, as well as those necessary for meaningful audits. Software integrity refers to the correct, unmodified software running on the electronic

components of the voting system. Faulty or malicious software may directly affect election data integrity.

5.2.1 Data Integrity Properties

Property: Accuracy

> *The election outcome properly reflects the choices of participating voters.*

Notes:

> The voting system must: (a) record votes consistent with voters' selections, (b) accurately store the collection of cast ballots, (c) protect the cast ballots from unauthorized modification, deletion or insertion, and (d) accurately count the votes.

Property: Auditability

> *The voting system provides evidence of its behavior before, during and after an election.*

Notes:

> It is not enough for a voting system to merely function correctly. The voting system must also provide evidence to auditors that the system functioned in the way it was supposed to. The evidence could include system event logs, public voting system reports, voter-verified records, and, in some cases, mathematical proofs. In addition, the voting system and its supporting election procedures must provide assurances that the evidence provided by the system is trustworthy. Auditability is a high-level security property of a voting system with more specific sub-properties listed and described in this sub-section.

Property: Privileged verifiability

> *The voting system provides evidence that allows the election auditors to independently check the outcome of the election.*

Notes:

> In general, verifiability is a voting system property where an observer is able to check the election outcome produced by the voting system is correct. That is, the system should produce ample evidence allowing auditors to verify the results of an election. In the case of privileged verifiability, the evidence provided could be secret or sensitive information that could only be made available to, or authenticated by, election insiders.

Property: Public verifiability

The voting system provides evidence that allows the general public to independently check the outcome of the election.

Notes:

Public verifiability is a property offered by emerging cryptographic voting protocols. In this case, sufficient evidence is made publically available by the voting system so any individual can verify the outcome of the election. Generally this requires some assumptions about the behavior of other entities (e.g., other voters, poll workers, administrators, etc.).

Property: Traceability

The voting system maintains all the necessary information so that if a problem is found in a particular election, then it is possible for the election officials to trace the problem to one or more root causes.

Notes:

If there are any problems during an election, it is important to be able to trace problems back to their root causes. The voting system should log or otherwise track sufficient events on the voting system to determine which activities failed or succeeded.

Property: Recoverability

The voting system maintains necessary information to allow it to recover from a loss of integrity. If the integrity of election records is lost in a way that is irrecoverable, then the extent to which the problem affects the final tally is measurable.

Notes:

If a voting system fails, then it should fail in a graceful manner. A minor problem should not necessarily call into question the integrity of the entire election. When possible, the voting system should be able to recover from minor problems. In some instances a voting system will not be able to recover from an error. In these instances, it should be possible to measure the extent of a failure so appropriate remediation can be carried out.

Property: Prevention of data alteration

The voting system prevents the unauthorized modification, deletion or insertion of election or voting system records.

Notes:

A voting system contains a great deal of data (e.g., system files, election records, and event logs) that must be protected from

unauthorized manipulation. To the extent possible, the voting system should prevent unauthorized manipulation and detect any manipulation that takes place.

Property: Logging data alteration

The voting system keeps a secure log with the information about who created/modified/deleted data which may influence the outcome of the election.

Notes:

Secure audit logs can help to increase accountability of system administrators and other insiders with privileged access to the machine. The log should be secure against modification by anyone, and should only be readable by authorized users.

Property: Data authenticity

Election auditors are able to verify the authenticity and provenance of election records.

Notes:

While protecting ballot secrecy, the voting system should provide sufficient evidence to allow election auditors to determine what entity (e.g., voter, system administrator, voting system component) created an election record and to verify that the record was not modified by unauthorized parties.

5.2.2 Software Integrity Properties

Property: Integrity of server software

Voting system components only load and execute authorized software.

Notes:

The voting system back-end components, such as servers, databases, and supporting network components, should only run authorized software. Front-end components under the control of election officials, such as kiosks, should also only run authorized software. For instance, the system should be free of malicious software. In addition, processes should be put in place to validate and authorize updates to voting system application software other third-party software used on the systems (e.g., operating systems, database software, anti-malware software).

Property: Authenticity of server software

Election auditors and/or system administrators are able to verify that only authorized software is present on voting system components.

Notes:

Auditors and system administrators should be able to verify that the voting system is free of malicious software and that only the authorized software is present on the voting system. In general, software validation is difficult to do rigorously and letting the voting system software verify itself is not sufficient.

Property: Proper software engineering practices

The voting system software is designed, implemented, tested and deployed with accepted software engineering best practices.

Notes:

Software engineering and testing best practices help to reduce errors in the design and implementation of voting systems.

5.3 Threats to Integrity

In general, any electronic system is prone to software bugs and malicious software attacks. Bugs and attacks related to software may result in partial loss of data integrity, and thus directly influence the election results. Moreover, Internet voting uses personally owned and operated devices which may be highly vulnerable to attacks that are capable of impacting election integrity. The election officials may have no practical way to assess the integrity of personal computers.

5.3.1 Software Bugs

One of the greatest threats to the integrity and accuracy of election records, including cast ballots, comes from non-malicious software defects, called software bugs. Software bugs accidently written into voting system application software, third-party libraries, and commercial software required to run the voting system all have the potential to impact election integrity.

Software bugs should be expected when dealing with software. In general, the larger a piece of software is, the more bugs are likely present. Estimates on the software industry's rate of bugs range from about 15 to 50 errors per 1000 lines of code [11]. Modern voting system application software can be quite large containing tens of thousands of lines of code. In most cases, voting systems run on top of commercial operating systems which can have

tens of millions of lines of code and use various other commercial libraries of software applications of varying complexity.

Extensive testing and analysis can identify many bugs but will never uncover all of them. Software bugs occur in medical devices, military equipment, and space exploration vehicles, despite extensive and sophisticated testing in these areas. In addition, software bugs affecting cast votes have been identified in certified voting systems [12], despite testing and code review during testing.

Even software whose source code is freely available to the public can contain major software bugs for years without discovery. The OpenSSL library included with the Debian-based linux distributions included a software bug in the cryptographic key generation function that resulted in a serious vulnerability in applications that relied on this library [13]. The bug went unnoticed for more than one year before being patched.

5.3.2 Malicious Software on System Servers

Specialized software could be maliciously placed on voting system equipment to modify or incorrectly store election records. The malicious code could be placed on the voting system equipment at any time in the system's life cycle. Developers of the voting system software, or any software used by the voting system, could include malicious code. Election insiders, such as system administrators, could install malicious software that changes election data. Or, remote attackers may be able to exploit a vulnerability in the voting system to install malicious code on the system. These attacks have the potential to change a large number of votes and can be difficult to detect.

5.3.3 Modification of Election Records and Data

Rather than installing malicious code on voting system servers and other back-end components, attackers may be able to modify election records directly to compromise election integrity. For example, a system administrator may be able to modify records stored on a database server. Or, vulnerabilities in the voting system may allow a remote attacker to perform an SQL injection attack to modify records in the database.

5.3.4 Malicious Software on Client Systems

The threats described in the previous sections are largely variations on similar threats faced by polling place electronic voting systems. However, Internet voting systems are also faced with threats to voters' personal

computers which are used as voting terminals. Attacks on these systems fall into a category generally referred to as client-side attacks. In most cases, these involve an attacker infecting a victim's computer with malicious software (e.g., a computer virus, trojan or worm) in order to gain access to information stored on that client machine or control it in various ways.

Client machines are quickly becoming a predominant attack vector in all types of information technology systems. Given the amount of sensitive information often stored on web, file, and database servers, these servers are often very tempting targets for attacks but they also tend to be the best protected, with professionally trained system administrators configuring and monitoring those systems. Client machines, used by regular employees or individuals, are often much less protected against attacks since they are operated by less technically sophisticated users. The client machine may be the intended target of an attacker, or it may be used as a stepping stone to attacking another computer system.

Attacks can use a variety of means to get malicious software on individuals' personal computers. Historically, file attachments sent over electronic mail were a common method for distributing malicious software. Alternatively, an attacker could post malicious software that appears to have a desirable purpose (e.g., a game, anti-virus software, screensaver, etc.) on a web site and encourage people to download it. In these cases, the victim became infected with the malicious software by executing the file attachment or downloaded file.

More recently, vulnerabilities in commonly used software became a common attack vector for malicious software. Some malicious software is self-replicating, where infected machines seek out other machines to infect, such as the 2003 Blaster worm that exploited a vulnerability in the Windows operating system. Individuals could become infected with the Blaster worm merely by connecting their Windows computer to the Internet. New vulnerabilities in commonly used application software have led to a new attack method, commonly referred to as drive-by-downloads. Using vulnerabilities in browsers, browser plugins, and other commonly used software, users can become infected with malicious software merely by visiting infected web sites.

An infected machine is largely under the control of an attacker. If a voter's personal computer becomes infected with a malicious software targeting the election, an attacker can potentially steal the victim's authentication credentials (e.g., a password or PIN), or even change the victim's vote without the victim noticing.

Malicious software is a serious problem on the Internet, with a large number of computers already infected with some type of virus or trojan. A growing problem on the Internet is botnets: groups of infected computers under the control of an attacker. The malicious software running on infected computers in a botnet is often used to steal passwords and other credentials for email and social networking sites in order to facilitate spreading the software to other computers. In many cases, the purpose of the attack is to steal financial data, such as passwords to online banking sites or credit card numbers. In some cases, malicious software on botnet-infected computers can even change the data inputted on a website for a financial transaction. For instance, it can change the bank account destination and amount for a money transfer on an online banking website.

Botnets are sometimes rented or sold by the individuals that originally conducted the attack to other parties. In addition, the malicious software behind the botnets is sold on black-market websites. For example, the malicious software behind the Zeus botnet is sold for as little as $700. Researchers at Cisco found that attackers could build a complete Zeus botnet for $2500, which includes the cost for the Zeus malware, exploit tools to infect users, and servers for conducting the attack [22]. While existing malicious software would have to be modified to attack an Internet voting system, this may not be difficult. In fact, many existing botnets include the ability to remotely update the malicious software running on already-infected computers. This means attackers would not necessarily have to re-infect computers already in botnets to attack an Internet voting system.

Because voters' personal computers are outside the control of election officials and voting system administrators, client-side attacks are very difficult to mitigate. While each successful attack on the client can only impact one vote or voter (or potentially a small number of voters if a computer is shared), attackers have demonstrated an ability to infect a large number of clients, and thus client-side attacks have the ability to have a large-scale impact.

5.4 Current and Emerging Technical Approaches

There are a number of techniques, some more mature than others, which can be used to address some of the threats to integrity of election results in the context of remote electronic voting. A summary of these techniques is presented below.

5.4.1 Cryptographic Integrity Protection

The data which travels through the Internet between the voting system and the voter's computer can be efficiently protected from en-route modifications via protocols like Secure Socket Layer (SSL) and Transport Layer Security (TLS). SSL and TLS are widely used to protect the integrity of communications between web servers and browsers and are frequently used in other applications as well such as email client and server communications. When used with mutual authentication, these protocols provide end-to-end security. In addition, cryptographic integrity protections, such as digital signatures and message authentication codes, can detect any changes in data as it is transmitted from one system to another. Cryptography can be very effective at protecting data in-transit and at-rest. However, cryptographic integrity protections do little to protect data as it is being processed on voting system components, such as when cast votes are constructed on client machines, or when they are tabulated on back-end equipment.

5.4.2 Advanced Cryptographic Voting Techniques

A specific research area in cryptography has been investigating more secure voting protocols to protect ballot secrecy, while at the same time offering unique integrity protections. These protocols, often called end-to-end cryptographic voting protocols, may be able to detect certain types of attacks where the election outcome is not the result of the aggregation of all cast votes. They can produce irrefutable proofs of tampering, even if a small number of cast ballots have been modified or deleted. Both voters and the general public can check that all cast ballots have been correctly tallied by the voting system. Additionally, end-to-end cryptographic voting protocols may allow each voter to verify that his/her vote appears in the final tally. There is a high degree of overlap between these protocols and the cryptographic protocols previously described in Section 4.4.2 to protect ballot secrecy.

The threat model for end-to-end cryptographic voting systems often assumes attackers have compromised the back-end voting system software. Thus, these systems can provide protection against attacks when cast ballots are modified in-transit or stored on voting system back-end equipment, and attacks that modify ballots or cause them to be incorrectly tabulated.

However, there are many types of attacks on voting systems that are not mitigated by end-to-end cryptographic voting protocols alone. In general, end-to-end cryptographic voting protocols may do little to mitigate client-side security threats, as cast ballots can be modified as they are constructed on the client machine. While end-to-end cryptographic voting protocols allow

31

the voter, or their proxy, to detect changes to cast ballots after they are constructed on a machine, they provide limited or difficult means to check the constructed cast ballot actually corresponds to the voter's choices. However, systems that provide clear text receipts of voters' choices are much easier to check, but these systems present potential problems with ballot secrecy and coercion. In addition, end-to-end cryptographic voting protocols do little to protect against attacks where voters' authentication credentials are stolen or sold.

At this time, there have been no formal usability or accessibility studies of current cryptographic voting schemes, but researchers anticipate that such studies would identify issues that would need to be addressed.

Remote electronic voting systems using end-to-end cryptographic voting protocols have been fielded in a limited number of pilots, including a university election at the Université Catholique de Louvain in March of 2009 [23]. End-to-end cryptographic voting protocols are an ongoing area of research, and researchers in academia and industry are coming up with different methods to address the shortcomings of these techniques.

5.4.3 Use of a Voter-Held Trusted Hardware Component

The threat posed by client-side vulnerabilities might be significantly reduced if the voter could use a third computing device that could communicate with the client machine and which could reasonably be assumed to be secure. Smart-cards and cell phones could, in principle, play this role. But it may be too expensive to add this capability to these devices for the sole purpose of voting, but this could be implemented to also help secure electronic commerce transactions.

5.4.4 Malware Detection/Prevention Software

Many commercial and free tools protect against malware, including antivirus and anti-spyware programs. These tools typically work by scanning files downloaded or opened on a computer. The tools look for patterns in files that match those of known malware. This is called signature-based detection. Many newer forms of anti-malware software can do more sophisticated heuristic-based checking in addition to signature detection to identify new malware. However, this is generally only effective at identifying new variants of an already-known piece of malware.

Anti-malware programs do not completely mitigate the threat of malware. Because anti-malware programs are dependent on an up-to-date list of malware signatures, users must update their anti-malware programs

frequently. In addition, anti-malware programs are not effective against new types of malware that have not yet been identified by vendors of anti-malware software and added to signature lists. Even known malware can be difficult to detect, as there are several techniques for writing malware to try to avoid signature-based detection. Once a computer is infected with malware, antivirus software may fail to detect or remove the virus. Some malware disables anti-malware software running on infected machines in ways that are not easy to detect.

5.4.5 Remote Software Verification

One area of research and development is remotely verifying that a piece of software on a given computer has not been tampered with. The most common application for this technology is to limit cheating in online gaming. In some online games, hackers have discovered ways of modifying software on their system to give them an unfair advantage. These anti-cheating mechanisms check the integrity of gaming software and data files looking for known cheating software in memory. It may be possible to extend these ideas to remotely inspect a voter's computer for malware.

Some current virtual private network software distributions include mechanisms to do end-point security scanning. When connecting to a server, the client machine downloads software from within the browser (often a Java application or ActiveX control) which performs some security scans on the client machine and relays the results to the server. Typically the purpose of scanning the system is to enforce an organization's security policy, such as running up-to-date antivirus software and a properly patched operating system.

An area of active research and development that may bring about more rigorous methods for remote software attestation is trusted computing platforms. In the future, it may be possible to use trusted computing modules (TPM) in voters' computers to demonstrate to an Internet voting system server the computers are in a desired state free of malware capable of tampering votes. The use of TPMs in voting systems is an active research area, with researchers proposing different methods for their use in voting systems [22][24]. While much of this research is focused on using TPMs in Direct Record Electronic (DRE) systems, the ideas could be extended for use in personal computers and Internet voting system servers. However, there are significant technical challenges to finding a workable solution. Furthermore, if and when solutions are found and implemented, deployment of the necessary hardware and software would likely be slow.

5.4.6 Formal Verification of Software

Formal verification of software involves providing mathematical proofs of the correctness of a given piece of software. In order to do formal verification, it must be possible to very precisely describe correct behavior in an algorithm. For this reason, formal verification is very hard to do for large software systems since it is difficult to precisely capture the behavior of a complex system. However formal verification is sometimes done for smaller pieces of a larger software system, such as the software implementing a cryptographic algorithm or protocol. Formal verification of software is very expensive, and is only done in extraordinary applications. For example, the INTEGRITY-178B real-time operating system, one of only two formally-verified operating systems, is used in military and commercial avionics.

Formal verification of system designs, while still uncommon, is required at Evaluation Assurance Levels 5, 6 and 7 of a Common Criteria security evaluation [25]. Again, these often involve verifying only a small piece of software within a larger system.

Because of its considerable cost, formal verification of software or designs is likely not well-suited to mitigating risks of software defects or vulnerabilities in remote electronic voting systems.

5.4.7 Preconfigured Bootable Environments

One method proposed for dealing with client-side security issues on voters' personal computers is to give voters a known-secure voting environment. This could be accomplished by distributed bootable media, such as CDs, DVDs, or USB drives that have been preconfigured with security hardening, and for connecting only to the Internet voting servers.

However, this approach has several significant disadvantages. One of the arguments for remote electronic voting has been the difficulty of distributing election materials to voters. Bootable media would likely have to be distributed by mail and would pose similar delivery challenges, such as obtaining up-to-date mailing address information for each voter. In addition, it would be very difficult to guarantee the bootable media would work on the vast majority of voters' personal computers. And, perhaps most significantly, it may be very hard for voters to identify legitimate bootable media from fraudulent media. Rather than serving to protect voters from malicious software, this could provide an avenue for attackers to distribute their own bootable media with malicious software preinstalled.

5.4.8 Virtualization Technologies

A possible way of bypassing some of the logistical problems of creating and distributing bootable media may be to use virtualization technology to run a clean voting environment in a virtual machine. That is, software running on a voter's computer could simulate a computer free of malware. This could alleviate some of the problems associated with bootable media including appropriate drivers and ensuring the default configuration would be compatible with a given user's network. Nonetheless, there are still significant logistical problems associated with attempting to securely distribute virtual machine images to voters. And, as was the case with bootable media, there remains the potential problem of voters using virtual machines pre-loaded with malicious software.

Generally, virtualization technology has been concerned with protecting the host operating system that is running the virtual machine software from any malicious or unreliable software running on the virtual machine's operating system. However, vendors of virtualization technology are beginning to implement systems that provide some protection against unauthorized modification of virtual machines by applications running on the host operating system. This is an important feature, as the reason for using these virtual machines is to protect voters from any malicious software running on their computers.

5.4.9 Secondary Communication Channels

While many of the technical approaches described above attempt to deal with the problem of malicious software on voters' computers by either detecting the malicious software or preventing its installation, another approach is to try to make voting from an infected computer reasonably safe. There are methods that attempt to do this using a secondary communication channel between the voter and the election authority that is independent from the voter's channel to the election authority such as the Internet through his or her personal computer. This second channel could be used when voters mark ballots to prevent malicious code from modifying votes in a directed way, or it could be used to confirm voters' selections.

In the first case, voters could be given individualized code sheets with unique random codes assigned to each candidate or choice on the ballot. In this case, the second channel might be the postal mail. To vote for a particular candidate, the voter would have to enter the random code assigned to that candidate on the Internet voting website. Malicious code running on the voter's computer would not know the association between the candidates and random codes, and thus would not be able to change votes to a particular candidate. However, malicious software could still

prevent voters from casting ballots, or try to deceive the voter into giving it the necessary information to change votes. In addition, there are significant usability concerns about this type of approach, in addition to logistical concerns involving the distribution of these code sheets to voters.

Alternatively, the second communication channel could be used to confirm a voter's selections. For example, a voter could be sent a message indicating how he or she voted. In this case, it is important that the second channel offer very fast delivery of messages, like a text message or telephone call, so the voter can confirm their selections in real-time. However, this approach creates some concerns related to vote selling by providing a channel which could be used by a vote buyer to verify how someone voted.

Electronic mail may be a tempting choice for a secondary communications channel, but there are significant drawbacks to using e-mail in this manner. E-mail is not an independent second channel, as the same computer and Internet connection used to construct and transmit the vote would likely be used to receive the e-mail. Malicious software running on the voter's computer may be able to change incoming e-mails along with cast votes.

5.4.10   Messages Computers Can't Understand

An alternative to using secondary communication channels is to communicate with the voter through the standard channel but coding information in ways that a computer cannot understand such as CAPTCHAs. CAPTCHAs are little puzzles that users are asked to solve, often involving reading distorted text, to prove that a human is accessing a Web application. CAPTCHAs are often used to try to block attacks where automated computer programs access a website and attempt to submit or collect information.

In principle, the whole ballot could be rendered using CAPTCHAS with the voter exercising choices by clicking on the rendered image. In this case, the client-side machine is unable to associate voter choices with locations of clicks. Even without the use of CAPTCHAS, using pointers to images instead of text should make it harder for malware to decode voter choices in order to alter them in favor of a given choice, because this is not a feature offered by currently available malware kits. Further research on these ideas is needed to identify usability and other issues that may arise. Note, these techniques do not stop the client machine from preventing the vote or randomizing it, and introduce usability and accessibility challenges that may not be adequately addressed.

5.5 Open Issues

Ensuring the security of personally-owned computers remains a very serious open issue. At this time, there is relatively little jurisdictions can do to ensure that voters' computers are free from malware capable of changing ballots cast from those machines. Attackers have demonstrated an ability to infect large numbers of machines with malicious software. Although in the case of UOCAVA voting, attackers would need to successfully target the relatively small percentage of individuals' in the world that are eligible to vote as overseas or military voters. While remote software verification, trusted computing modules, and computer virtualization are potentially promising technologies for mitigating the threat of malware on voters' computers, none of these technologies appear ready for immediate use with remote electronic voting systems.

There are also open issues related to the security of software on voting system servers. While extensive testing may be able to uncover many software bugs, there are no guarantees it can uncover all bugs in the software.

Advanced cryptographic voting technique, specifically end-to-end cryptographic voting protocols, can be highly effective at detecting certain types of attacks on election integrity, including modification or deletion of cast ballots. However at this time, they are most effective against mitigating attacks that take place on the voting system servers. Most of these techniques are not effective at detecting attacks taking place on the computers used to cast ballots. While extending end-to-end cryptographic voting protocols to detect client-side attacks is an active research area, methods that have been proposed are either difficult to use or impractical. In some cases, end-to-end cryptographic voting techniques only detect if an integrity violation has occurred. It may not be possible to recover from the detected error or to measure the extent to which the detected error affects the outcome of the election. Also, end-to-end cryptographic voting techniques may not be able to distinguish between a bug and an active attack. While this is an area of ongoing research and activities, end-to-end cryptographic voting techniques for Internet voting are largely still an academic effort.

6 Availability

Availability is used to describe the proportion of time a system is functioning and operating, including times when the system is performing at reduced capacity. Due to resource overload, malicious attack, and system malfunction, a system may become unable to function, and thus is considered unavailable.

6.1 Potential Benefits

Electronic transmission of election materials can provide several benefits to UOCAVA voters and election officials compared to alternative voting methods for overseas and military voters. The following section describes some of the potential benefits.

6.1.1 Timeliness of Delivery

Internet voting systems do not suffer from the same delays associated with voting through the postal mail. Postal mail delivery to remote locations can take significantly more time than delivery times within the United States. For example, delivery through the military postal system to Middle East postal offices typically takes 7-12 days [27]. Internet transmission, however, is nearly instantaneous, as long as voting system endpoints (the server and the client) and communication lines are operational.

6.1.2 Receipt Confirmation

The United States Postal Service (USPS) is a relatively reliable delivery mechanism, with first class mail on-time performance exceeding 90% [28]. However, mail to UOCAVA voters must go through other postal services in addition to the USPS, such as the military postal system, or those of foreign nations. Delivery confirmation is an option for USPS mail to military addresses, but is often not an option for mail to and from foreign addresses. Therefore, it is nearly impossible to detect which blank or completed ballots have been lost or delayed in the mail system.

Remote voting over the Internet can provide immediate feedback to senders if there is a transmission problem via real-time confirmation and error messages. This information could be used to detect problems and remediate them.

6.1.3 Flexibility of Physical Locations

Overseas voters, particularly military voters, are a highly mobile population, and are not always quick to inform their local election officials of their new addresses. Remote voting over the Internet allows voters to receive or cast ballots regardless of their physical location.

6.2 Properties

Property: Up-time

Voters, election officials, and other system operators are able to use the voting system normally for a substantial percentage of the total time allowed to configure the system, cast votes, and tally votes.

Notes:

Up-time is a measure of the extent a system is available for use by system operators and users. A number of factors affect up-time, including how often failures occur (see the "Reliability" property) and time it takes system administrators to restore functionality after a failure occurs. System availability can be maliciously targeted by an attacker to disrupt voters from casting their ballots.

Property: Reliability

The voting system, to a high degree of probability, will remain operational during the election under predefined normal operating conditions.

Notes:

Reliability is a measure of the likelihood a system will continue to perform as intended for a specified time under a particular set of predefined conditions. In this case, reliability is referred to as the likelihood the voting system can complete an election without a loss of functionality when it is not facing a malicious attacker.

Property: Recoverability

Voting system operators are able to restore the system to normal operation in a timely manner when failures occur.

Notes:

Voting systems should be designed to limit downtime in the event of failures. In practice this implies a very low probability of catastrophic failure such as loss of stored cast ballots.

Property: Fault-Tolerance

The voting system is able to continue operation, perhaps at a reduced level of functionality, when failures or attacks occur.

Notes:

A common method for achieving some level of fault-tolerance is to use redundant system components or resources.

Property: Fail-Safe

In the event of a failure or attack, the voting system experiences minimal data loss or further damage to voting system components not directly affected by the failure or attack.

Notes:

Fail-safe is a system property which states that voting system failures or attacks should have limited impact on the integrity and availability of system components and data. For example, hardware component failures in the voting system should not result in the loss of cast vote records or audit information. An attack on one component in a voting system should not damage a second component. For instance, an attack on the voter registration database should not harm the voting system web server, although it may inhibit voting activities until the issue with the voter registration database is resolved.

Property: Scalability

The capacity of the voting system can be increased with additional resources (e.g., servers, network bandwidth, etc.) without redesigning the system's architecture.

Notes:

A scalable voting system can grow to accept greater and greater number of voters by adding additional hardware, more powerful hardware, faster network connections, other computing resources, or any combination thereof.

6.3 Threats to Availability

Like any information technology system, Internet voting may be the target of denial of service attacks (see [29] for precinct voting denial of service attacks). The potential scale and impact of the attack may be much larger for Internet voting systems than for polling place voting or mail-in voting. The attacks can be targeted towards the server providing the voting service, the personal computers of the voters, or the infrastructure connecting the two. Denial of service attacks may be selective, such as disrupting service for voters deemed likely to cast a ballot in a particular way (e.g., a particular demographic group).

6.3.1 Large-Scale Denial of Service

Denial of service attacks are a type of attack where malicious individuals attempt to make a computer system unavailable to its users. Depending on

the nature of the attack, and on its target, a denial of service attack can be anything from a minor nuisance to a devastating attack.

Denial of service attacks could prevent voters from being able to cast votes either by making Internet voting system servers inaccessible or disrupting systems they rely on, such as the communications infrastructure or voter registration database. Aimed at the back-end of the voting system, these attacks could prevent large numbers of people from casting ballots over the duration (anywhere from hours to days) of the attack.

Denial of service attacks of varying severity occur frequently on the Internet. The type of target and motivation differs from attack to attack. A frequent motive of attackers is political in nature, with attacks carried out by individuals or groups disagreeing with the victim's views. Large corporations, nation states, and the communications infrastructure are frequent targets for attack. For example, in 2007 the nation of Estonia was targeted with a large-scale denial of service attack [30], with the nation of Georgia experiencing a similar attack in 2008 [31]. Critical portions of the Domain Name System (DNS) have also been targeted with attacks, including distributed denial of service attacks against root DNS servers in 2002 [32] and 2007 [33].

Denial of service attacks can be conducted in a variety of ways, but most major attacks are distributed denial of service attacks. Collections of malware infected computers, known as botnets, can be purchased or rented by attackers to be used to attack a target organization.

6.3.2 Selective Disruption and Suppression

While denial of service attacks can cause voter disenfranchisement on a significant scale, their ability to impact the outcome of an election is somewhat limited unless the attack is focused on a particular demographic or jurisdiction. However, targeted denial of service attacks have been documented. In 2009, denial of service attacks targeted a specific Georgian blogger on Twitter, Facebook, Livejournal and Google [34]. Denial of service attacks that selectively disrupt systems at a particular jurisdiction or certain voter demographic could not only result in voter disenfranchisement, but also sway the results of an election.

Remote electronic voting may make it harder to prevent a voter from attempting to vote when the voting system is architected to function and operate even under vote suppression attacks. On the other hand, some cyber attacks, such as denial of service attacks, may make it easier to thwart an attempt to vote due to the resources available to an attacker in the form of computers controlled by botnets.

6.3.3 Client-Side Disruption

While most large-scale attacks on availability target one of the voting system's servers or the communications infrastructure, attacks can also target the voters' machines. Malware on voters' computers could prevent them from accessing voting web sites.

6.4 Current and Emerging Technical Approaches

While there is no general solution to denial of service attacks, a series of techniques can be used to prevent, detect and speed up recovery from such attacks.

6.4.1 Redundancy and Over-provisioning

The most widely used approaches for achieving high-availability systems include the use of redundant systems and over-provisioning of system resources. At a basic level, these approaches involve fielding systems with excess capacity so they are able to better handle failures on certain system components or attacks.

Redundancy involves the duplication of critical system components. The duplicate components are used as backups in the event of failures or to augment capacity in the event of a spike in legitimate or illegitimate traffic. For instance, a system could be designed with redundant web servers such that the backup system can take over the expected load in the event the primary system fails.

A more general approach, called over-provisioning of system resources, involves fielding systems capable of handling a much greater load than would be expected under normal conditions. A useful strategy is to identify possible performance bottlenecks in the system and to augment the capacity at those bottlenecks. Possible bottlenecks include capacity and performance of the communications lines, support infrastructure (such as firewalls and routers), or database and web servers. Over-provisioning can involve any combination of duplicating resources (e.g., mirrored sites located at multiple physical locations) or making individual resources more powerful or abundant (e.g., faster network connections, more powerful servers, etc.) than would ordinarily be needed.

Fielding over-provisioned systems can be costly, particularly for relatively small systems such as Internet voting systems that are rarely used and have less traffic than major e-commerce web sites. Small increases in system capacity are not likely to deter or prevent attacks on availability, but large

increases in capacity may be wasteful and still potentially ineffective. Over-provisioning raises the bar for attacks but does not make attacks impossible.

6.4.2 Detecting Active Attacks on Availability

Compared to other types of attacks on voting system, availability attacks are usually relatively easy to detect by system administrators. In some cases, the system crashes or becomes unavailable to all users. At this point, voters have already been affected and will continue to be affected until the attack is successfully repelled. The key to maximizing availability in the face of denial of service attacks is early detection and quick reaction.

6.4.3 Defending Against Active Attacks

The most common approach for defending against denial of service attacks is over-provisioning, which provides protection against all kinds of incidental or malicious threats to availability. However, there are a number of other things system designers and administrators can do to defend against attacks.

One approach is to preemptively harden systems against denial of service attacks. Hardening voting systems include identifying and fixing bottlenecks as well as vulnerabilities in host systems that make denial of service attacks easier to carry out, and carefully designing the internal network infrastructure. In some cases, there may be multiple technical options for designing a secure and usable voting system that works equally well for their intended tasks but may be more resistant to denial of service attacks.

Another approach is to filter dangerous network traffic containing known attacks carefully constructed to crash or overwhelm a particular system resource. Once an active denial of service attack is detected, an organization may be able to filter out the network traffic making up the attack. While network traffic filtering can be done on the border of an organization's network, an attack may attempt to overwhelm the filtering mechanism or merely fill the in-bound network connection. In these cases, it is helpful to filter attack traffic closer to the source, which usually requires the help of third-party Internet service providers.

Some distributed denial of service attacks work on the premise an attacking client can force a server or other device to consume far more resources than those required by the client to conduct the attack. For example, establishing a Transmission Control Protocol (TCP) connection with the server requires that the server allocate resources before the client. There are approaches that attempt to address the client server resource imbalance, such as SYN

Cookies and proof-of-work techniques, by forcing clients to allocate some resources before establishing a connection with a server [35].

6.4.4 Cloud Computing

In protecting system availability, there is strength in numbers. Having redundant systems to migrate to after a failure, or having excess capacity to raise the bar for denial of service attacks, can help systems achieve higher levels of availability. However, purchasing, deploying and maintaining this excess capacity may be cost-prohibitive. An emerging area in the computer industry is a concept known as cloud computing. Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction [36].

In the cloud computing model, a large pool of resources can be distributed between many different applications and even customers. Excess capacity in the system can be applied to any of the applications running in the cloud on an as-needed basis, and the cost associated with maintaining the excess capacity can essentially be distributed across all of the customers. In the event of a hardware failure on a particular machine in the cloud, any applications running on that machine can be almost immediately transferred to a different physical machine in the cloud. In the event of a spike in traffic for a particular application, additional physical or logical machines, network bandwidth, or other resources could be allocated to that application.

However, in cloud environments, multiple applications are being hosted on the same systems. So, in the case of an Internet voting system, the voting system may be running on the same equipment used to perform completely unrelated tasks. When a service provider manages the cloud, each customer may have little control of what other applications coexist on the same physical equipment. Typically, virtualization technology is used to keep different application resources logically, rather than physically, separate. However, this introduces new security issues researchers have only begun to look at in the last few years.

Cloud computing appears to be a very promising technology for increasing system availability in a cost-effective manner, but it is not clear if it is ready or suitable for use with remote electronic voting systems.

6.5 Open Issues

Most defensive techniques against denial of service attacks can raise the bar for an attacker to successfully mount an attack but cannot guarantee protection. In fact, due to the nature of the Internet, it may not be possible to provide complete protection from certain types of availability attacks. Given the commercial availability of botnets for use in distributed denial of service attacks, attacks on availability are a very real threat to Internet voting systems.

However, Internet voting systems are no more vulnerable to denial of service attacks than many other types of online computer systems as, at a high-level, their architectures have many similarities. And, the threats to voting system availability should be considered relative to availability issues faced by mail-in absentee voting, including undeliverable mail due to a frequently moving overseas voting population and the time necessary to send or return election materials.

Cloud computing appears to be a promising technology. However, it is a young field where researchers and developers in industry and academia are making advances at a rapid pace. The security issues associated with cloud computing, along with new types of potential vulnerabilities, continue to be identified.

7 Identification and Authentication

Determining if a user is authorized to use a voting system includes the distinct steps of identification and authentication. Identification is the act or process in which an entity (e.g., user or system component) provides a unique identity so a system can distinguish the entity from all others. Authentication is the process of establishing confidence in user identities.

Proper voter authentication is required to ensure only eligible voters can cast ballots and a valid voter contributes a single ballot to the final tally. A remote voting system will typically verify credentials it is provided with, and assume the person providing those credentials is the legitimate owner. As credentials may come from the voter's computer rather than from the human voter him or herself, the voter's computer may gain direct, unrestricted access to the voting credentials. The binding between voters and identities, and between identities and credentials, is established through "voter identification."

It is also important, in a remote setting, that the voting system authenticates itself to the voter. This implies that the voter is able to check that she is actually interacting with the legitimate Internet voting service.

7.1 Potential Benefits

Polling place voting typically authenticates voters by having polling place officials interact directly with the human voters. In some cases, voters may be asked for identification or some other authenticator. In Internet voting, strong cryptographic credentials can be used to authenticate voters. In such cases, cryptographic authentication mechanisms make it essentially impossible to trick the system into accepting forged credentials.

7.1.1 Automated Authentication Mechanisms

Hand signature verification generally requires trained election workers to inspect every ballot package returned by voters, matching the signature included with the ballot to a signature specimen on file. While some absentee voting management systems can automate some of the signature comparisons, it is still a moderately resource intensive activity. However, electronic authentication methods can be entirely automated.

7.1.2 Strong Remote Authentication

Currently, remote electronic authentication methods exist which are capable of providing high levels of assurance of a user's claimed identity. Many of these methods are widely deployed in the public and private sectors.

Although the stronger authentication mechanisms are typically used in government, military or corporate environments, they have not been widely deployed to general public. For instance, the federal government's Personal Identity Verification program of the federal government involves distributing smart cards to government employees and contractors for authentication purposes. The Department of Defense's Common Access Card is similar program for military personnel and contractors. However, most citizens of the United States that are not associated with the federal government or military so do not have smart cards. The situation in the United State is different from other countries that have deployed Internet voting systems, such as Estonia, which have smart cards deployed to the vast majority of the general population.

7.2 Properties

Property: Voter Identification

Election authorities and voting systems are able to uniquely identify eligible/registered voters within a particular jurisdiction.

Notes:

Unique identification of voters is necessary to bind eligible voters to digital identities and digital identities with credentials. The credentials are used for voter authentication and enforcing access control rules and keeping records of who did what on the voting system.

Property: Voter Authentication

The voting system verifies the credentials of potential voters before allowing them to perform any authorized actions on the system.

Notes:

The voting system should ensure that voters connecting to the system are eligible to use the system to perform the requested functions (e.g., cast a ballot, update voter registration information). In remote authentication, it is important to understand there is no difference between authentication of voters and authentication of credentials. That is, anybody with access to the voter's credentials is able to impersonate the voter. There is a spectrum of techniques that offer different levels of assurance in remote authentication. For example, 4-digit pins offer lower remote authentication assurance than strong passwords. Higher assurance can be obtained using "two-factor authentication" methods typically involving cryptographic token and a PIN, a password and a biometric, or a time-dependent random number generated by a small hardware device issued to the user. Voting system authentication in the foreseeable future is unlikely to make use

of biometrics, but deployment of some form of two-factor authentication does seem feasible for special populations such as military personnel. For instance, the Department of Defense has distributed the smart card-based Common Access Card (CAC) [43] to nearly all of its military personnel, employees and contractors.

Voter authentication should not compromise the secrecy of the vote. The authentication protocols should not attach easily retrievable or inferable voter identification information to cast ballots. If jurisdictions allow a voting system to attach voter identification information to cast ballots, then this information should be encrypted in such a way that it can only be decrypted under exceptional circumstances.

Property: Administrators/Officials Identification

Election authorities, system administrators, or other individuals with administrative access to voting systems, are uniquely identified by the voting system.

Notes:

Individuals with privileged access to the voting system should be uniquely identified by the voting system. That is, system administrators, election officials, and other with access to voting system should not share accounts or login credentials. This allows for greater accountability of administrative actions performed on the voting system.

Property: Administrators/Officials Authentication

The voting system components verify the credentials of system administrators, election officials, and other election insiders before allowing them to perform any actions, as authorized, on the system components.

Notes:

Voting system administrators and election officials do not require the same privacy protections as voters. Thus, every voting system component should verify the unique identity of the official or administrator before granting them access to the system.

Property: System Component Identification

Each voting system component is identified by the system.

Notes:

Like users, each voting system component should be identified. While some level of unique identification would be necessarily for various

administrative functions for logistical reasons, groups of components that act as one might be identified as part of a collective group. For instance, individual machines in a group of web servers behind a load balancer may all share the same identity for identification and authentication purposes.

Property: System Component Authentication

Users and system components should verify the identities of voting system components before any other interactions with those components.

Notes:

It is important to note that this property applies both to users (e.g., voters, election officials, administrators) connecting to voting system components, as well as voting system components connecting to other components. In both cases, users and voting system components connecting to the voting system should verify they are communicating with the component they intended and not some other computer system impersonating the intended component. In particular, voters should authenticate the voting system they are interacting with, to ensure it is the legitimate voting system.

Property: Non-transferable Credentials

It should be difficult for voting system credential holder to pass his or her credentials to an unauthorized party without detection.

Notes:

Section 7.3 discusses several threats to identification and authentication systems where an attacker convinces a legitimate user to disclose credentials to an unauthorized party. In most cases, this would involve deceiving the legitimate credential holder but could be done with the cooperation of the credential holder (e.g., in the case of vote selling). Credential transfer attacks should be difficult to perform without detection. In this case, difficult may mean the attack does not scale well, or that the threat of punishment if caught is severe enough to deter attacks.

7.3 Threats to Identification and Authentication

7.3.1 Unauthorized Issuance of Credentials

One common threat to identification and authentication systems is that unauthorized parties may be issued credentials they are not eligible for. For instance, an individual may impersonate some other individual and register

in his or her name. Alternatively, an individual who is not eligible to vote in a jurisdiction may register to vote and be issued credentials to vote. These types of threats are very similar to current forms of voter registration fraud.

There continues to be disagreement over the extent and severity of voter registration fraud in the United States. A study of election crimes by the Election Assistance Commission found that while experts agree fraudulent voter registration forms are filled out, most do not believe these fraudulent registrations result in fraudulent votes actually being cast [37].

It is not known how a move to remote electronic voting over the Internet will change the threat environment for these forms of voter registration fraud.

7.3.2 Phishing/Pharming

Phishing and pharming are two related attacks on the Internet today. While the method for conducting the attack differs between the two, the goal of the attacker is the same: to trick users into revealing their credentials on an illegitimate web site that looks like the legitimate site. In the case of phishing, an attacker sets up a fake website and lures users to the site. Attackers use a variety of means to lure users to these websites, but they typically involve registering a website domain name similar to the legitimate web site and sending mass e-mails claiming to be the legitimate website owner but including links to the fake website. Phishing is largely an attack on the user, rather than on any particular piece of equipment. Pharming is a similar attack, except rather than tricking a user into visiting the fake web site, attackers use some sort of computer or network vulnerability to redirect a user from the legitimate website to an illegitimate one without the user's knowledge.

Phishing attacks are very widespread on the Internet, with credentials to financial and social networking sites often being the target of the attacks. According to a Gartner report, five million consumers in the United States lost money to phishing attacks in fiscal year 2008 [38]. Their survey estimated the average consumer loss per successful phishing attack was $351. However, accurate information on the losses associated with phishing is very difficult to collect, and other researchers have questioned the accuracy of this information, claiming that actual losses are much lower [39]. A recent report by the Anti-Phishing Working Group found phishing attacks continue to be a significant problem, with a record number of organizations targeted by phishing attacks in the fourth quarter of 2009 [40].

Phishing and pharming attacks on Internet voting systems could successfully steal voters' credentials, allowing malicious parties to cast votes in place of the legitimate voters. Attackers may also conduct more targeted phishing attacks, sometimes called spear phishing, on election system administrators or election officials, possibly resulting in gaining privileged access to back-end voting system equipment. Because these attacks are just as much attacks on human users as they are on the technical system, they are very difficult to prevent. Phishing attacks in particular require very little resources and technical expertise to conduct, yet can impact a very large number of people. While Section 7.4.5 will discuss a common method for preventing phishing and pharming attacks, its benefits are somewhat limited.

7.3.3 Credential Selling

Some types of credentials are very easy to transfer to another individual. For instance, PINs and passwords can be physically or electronically sent to another individual as part of a vote selling attack, as described in Section 4.3.3 or in attempts to coerce voters, as described in Section 4.3.2. As noted in those sections, it is difficult to estimate the likelihood of such attacks or how motivated potentials attackers would be to conduct these types of attacks. However, depending on the types of credentials used, these attacks could scale fairly well, potentially allowing individuals or organizations to collect large numbers of voters' credentials and cast votes on their behalf.

There are technical measures that could be taken to greatly limit the ability of these attacks to scale, such as using credentials that cannot be easily passed from a voter to another individual. For instance, use of hardware tokens, such as smart cards or one-time password devices, could require a voter and coercer/vote-buyer to exchange a physical device. However, these mechanisms typically come at a higher cost than simple authentications based on passwords or PINs. Biometric characteristics used in conjunction with challenge-response protocols may also be used to make it impossible to transfer a person's credentials to someone else.

7.3.4 Social Engineering

Social engineering is a class of attack where malicious (or curious) individuals manipulate legitimate users of a system into divulging sensitive information, such as login credentials for a system. Phishing and pharming can be considered a type of large-scale, automated social engineering attack, but social engineering attacks could be highly targeted and interactive. For instance, an attacker conducting a social engineering attack could call an election official or system administrator claiming to be from the service provider hosting the voting system and convince the victim to divulge his or her password.

Social engineering is a class of attacks, and the objective of the attacker may not be solely to steal login credentials. The objectives of social engineers can be to obtain any type of sensitive information that may help them conduct an attack.

7.3.5 Cracking/Guessing

Depending on the type of authentication mechanism used and the location of the attacker, a malicious individual may be able to steal authentication credentials with brute force. This is particularly true for authentication mechanisms like passwords or PINs, as well as knowledge-based authentication. For example, a randomly-generated four-digit PIN has ten thousand different possible values, so an attacker has about a 0.5% chance of guessing a PIN after 5 attempts. In the case of user-chosen passwords, people tend to choose dictionary words for passwords, making it easier for attackers to guess or crack a password.

There are a number of methods that system designers can use that can make it very difficult to guess or crack a particular individual's login credentials. However, if a system has a large number of users, it is much more difficult to ensure that none of the users' credentials are cracked or guessed. This may not be a serious concern for voters' credentials, as these attacks do not appear to scale well.

More seriously, individuals with some level of access to the system, such as physical access to voting system equipment or the ability to watch network traffic between voting system components, may be able to use more sophisticated cracking or guessing attacks. This could be the first stage of an attack if the person is some sort of election system insider (e.g., a computer technician at the service provider hosting the system), or it may be done by a remote attacker that has already gained limited access to the voting system equipment. The impact of these attacks can vary. An attacker that successfully guesses or cracks the credentials associated with a privileged account would be able to perform any actions on the system as if they were the legitimate user.

7.3.6 Malicious Software

Malicious software, or malware, on computers of users' connecting to the voting system could steal credentials used to authenticate to the system. For instance, a common example of malware used by attackers is a keylogger. Keyloggers can record everything that users type on their keyboards. Therefore, it is capable of capturing authentication credentials like

passwords and PINs very easily and can pass them to a remote attacker over an Internet connection. Keylogging functionality is common in malicious code associated with botnets, which were previously discussed in Section 5.3.4.

As was the case with credential guessing and cracking, the impact of these attacks can vary. Attackers that steal the credentials associated with a voter's or administrator's account would be able to perform any actions on the system as if they were the legitimate user. This means that attackers may be able to cast votes in place of a voter, or even perform administrative functions if they are able to get malicious software on a computer used for system or election administration.

7.3.7 Insiders/Credential Issuers

If voting credentials are issued by a particular entity, such as the election officials giving voters usernames and passwords, these insiders have access to all the credentials used for casting ballots. Such an individual may use these credentials to cast votes in the name of voters (for example for voters who did not cast ballots until a couple of minutes before the polls close).

To avoid such scenarios, it may be best to have the voter choosing their own credentials, with insiders never having access to these credentials in clear text, but at the same time being able to check that the voter have knowledge/access to them. For example, if electronic signatures are produced using smartcards, the private keys have to be generated inside the smartcards and it should be impossible to read the clear text private keys, but only to use it to sign messages.

7.4 Current and Emerging Technical Approaches

7.4.1 Passwords and PINs

Passwords and PINs remain two of the most common methods for electronic authentication, largely because they are relatively cheap and easy to deploy. Most people use passwords to log into their computers and web-based accounts, including e-mail, social networking sites, and financial sites. Passwords and PINs are typically user-generated, although in some cases organizations or systems will send users pre-generated passwords initially and ask the users to change them when they are first used.

However, passwords and PINs have significant security disadvantages compared to other types of authentication mechanisms. User-generated passwords can often be easily cracked if the attackers have sufficient

information, and they are easily stolen by malware or phishing sites. For these reasons, many organizations are moving away from just using passwords for authentication. For instance, the federal government requires some form of two-factor authentication for remote access to government systems [41], and some financial institutions have begun using two-factor authentication for online banking.

7.4.2 One-time Passwords

One-time passwords are a common method for deploying two-factor authentication. A one-time password is a password that is only valid for a single transaction and usually a short period of time. In most cases, systems using one-time passwords still use user-generated, memorized passwords, with the one-time password adding an additional layer of authentication.

The difficulty of one-time passwords is organizations need a method for securely distributing these one-time use passwords to their users. This is typically done one of two ways: distributing trusted hardware devices to users or sending them on-demand through a secondary channel such as a cell phone.

Many organizations in the public and private sectors use trusted hardware devices to generate one-time passwords. Organizations must keep track of which users are given what one-time password device. These devices typically continuously generate random codes at regular intervals, such as every 30 seconds. When a user attempts to log into a system, he or she typically must enter both a memorized password in addition to the random code on the one-time password device at that particular moment. The use of a hardware device increases the cost of the system, and the device must be securely distributed to users either in-person, or by some other physical means, and may be lost by users.

Alternatively, one-time passwords can be sent or generated on devices that users already have. For instance, a user may have a piece of software on his or her mobile phone that generates one-time passwords in a similar manner as the hardware device described above. Or an organization may have the mobile phone number for a user and send one-time passwords as text messages on-demand to users attempting to authenticate to the system.

The use of one-time password devices can provide some protection against the threats described in Section 7.3 with some important limitations. Because these passwords are constantly-changing strings, they are very difficult to guess or crack, so malware and phishing sites cannot easily collect large numbers of passwords for later use. However, more

sophisticated attacks can be conducted by malware and phishing sites. If an attacker can capture a one-time password and use it before the user sends it to the legitimate system, the attacker can successfully impersonate that user. This can be accomplished with phishing websites that immediately connect to the legitimate website when a victim enters his or her information, or with malware that passes credentials to an attack in real-time. Both of these types of attacks have been used to attack online banking sites and do not require particularly high-levels of technical expertise. Some malware packages commercially available, as were discussed in Section 5.3.4, include the ability the conduct these types of attacks [42].

7.4.3 Cryptographic Authentication

There are various forms of cryptographic authentication that can be done remotely using cryptographic tokens. These tokens are used in a cryptographic protocol whereby the user proves to the organization authenticating them that he or she has possession of the cryptographic token without having to directly present the token. Authenticating using cryptographic tokens can have very strong security properties and can be implemented such that they are very difficult to crack or steal via phishing.

Cryptographic tokens can be software or hardware based. The difference is whether the cryptographic token is stored on a trusted hardware device, such as a smart card, or whether it is merely a file or piece of software on a computer, mobile phone, tablet PC, or other general-purpose computing device. Software-based cryptographic tokens are vulnerable to theft or tampering but do not require any special hardware. Hardware-based tokens provide greater security.

Hardware based cryptographic tokens often take the form of a smart card. Smart cards are used by many organizations in the public and private sectors for authentication purposes. The Department of Defense has distributed the smart card-based Common Access Card (CAC) [43] to nearly all of its military personnel, employees and contractors. The United States federal government is in the process of implementing a similar program, the Personal Identity Verification card [44], for civilian employees and contractors. In lieu of issuing credentials specifically for voting, UOCAVA voting systems should consider leveraging strong credentials that are already deployed. For example, the country of Estonia, which has a smart card-based national identification card, performed voter authentication in its Internet voting system using the electronic credentials found on the national identification card [45].

Cryptographic authentication is also well-suited for allowing components of the voting system to authenticate to one another. There are a number of networking protocols that allow one component to authenticate to other components. Transport Layer Security (TLS), for example, is a commonly used protocol on the Internet to encrypt traffic between a website and a user's computer and to authenticate the website to the user's system. TLS can also be used to authenticate the client connecting to a server. While client authentication is relatively uncommon in typical e-commerce transactions, it is often used in higher security systems.

7.4.4 Biometrics

Biometrics are methods for identifying and authenticating individuals based on one or more behavioral or physical traits. Commonly cited biometrics used for authentication purposes include fingerprints, iris recognition, and hand/palm geometry. Biometric authentication can offer high degrees of security depending on the quality of the biometric readers used in the system. However, biometrics are typically used for local authenticating, meaning the user authenticating to a system is in the same physical location as the system. This is because biometrics must be measured by a trusted reader, such as a fingerprint scanner.

Some biometrics are better suited for remote authentication, such as speaker verification. Speaker verification authenticates a user based on their speech patterns. This should not be confused with speech recognition, which recognizes the spoken words, regardless of the identity of the person speaking. Currently, speaker verification methods provide significantly higher error rates than other biometrics [48], but it is an active research area with a number of commercially-available systems. Speaker verification may be suitable as a secondary authentication method or, with improvements to technology, a primary method.

7.4.5 Phishing Filters

Many modern web browsers and anti-malware software distributions include some type of protection against phishing attacks. These approaches typically involve some combination of whitelisting websites known to be safe, blacklisting websites known to be fraudulent, and, in some cases, using heuristics for all other websites in an attempt to estimate the risk of phishing (e.g., a URL using an IP address instead of a domain name). When a user visits a website that is deemed unsafe, the phishing filter displays either a passive or active warning. An example of a passive warning in a browser is a short warning message, such as "Suspicious Website," placed next to the address bar, but does not require any user input to ignore. An active

warning interrupts the user and requires some sort of input by the user to ignore the warning. For instance, before displaying the phishing website, a browser may display a warning page telling the user the website is a suspected phishing site and asking the user if he or she would like to proceed to the page anyway.

However, the effectiveness of phishing filters is limited by their ability to identify fraudulent websites and how well users heed the warnings. A 2006 study by the Mozilla Project found that between 66% and 82% of fraudulent web sites were detected by the phishing filters used in two popular web browsers [46]. A limited number of usability studies have been done on phishing filters. A 2008 study found that 90% of Internet Explorer 7 users ignored passive warnings from the browser's phishing filter, with that percentage improving to 45% when an active warning was displayed [47]. However, new designs for phishing warnings may be able to improve those rates.

7.4.6 Security Awareness and Training

Many of the threats described in this section are attacks on users, rather than on the voting system components. In some cases, users are not aware of the security threats faced by a system, or what actions might pose a security risk. Security awareness presentations and materials can educate users about these threats in the hopes that they will be less likely to fall to a phishing or social engineering attack and more likely to use safe computing behaviors. Security training can educate users about relevant security skills and competencies that are necessary for them to conduct their jobs effectively and safely.

Jurisdictions should develop security awareness and training programs for election staff. They may also distribute security awareness materials to voters highlighting recommended security practices and potential threats.

7.5 Open issues

Unlike some the other topics areas described in this document, many of the security challenges associated with identification and authentication of users and voters have commercially-available technical solutions. However, there remain logistical concerns, as well as concerns over the cost of implementing some of these solutions.

Deployment of strong authentication credentials for voters is an issue that would likely be difficult for jurisdictions to manage at this time and could be difficult for the foreseeable future. The authentication methods providing the

highest levels of assurance of users' identities involve specialized hardware devices that increase the cost of the system and complicate deployment. It may be advantageous for jurisdictions to rely on already deployed authentication credentials, such as the DoD's Common Access Card and the federal government's Personal Identity Verification card, which are already deployed to many overseas voters. However, it is not known if these credentials could be used for voter authentication, or what would be done with the hundreds of thousands, if not millions, of overseas voters that do not have one of these electronic credentials. This could change over time; as more people conduct electronic transactions in their daily lives, it may become increasingly important for all citizens to have strong electronic credentials.

The threat of phishing and social engineering attacks are logistically, and even technically, difficult to mitigate. Cryptographic tokens can provide some protection against phishing attacks, but many other authentication techniques can still fall to variations of phishing and social engineering attacks. Mitigating phishing attacks will likely require a combination of technical controls, possibly in the form of cryptographic tokens, and users better able to understand risks and identify risky behavior.

8 Conclusions

This paper identified desirable security properties of remote electronic voting systems, threats of voting over the Internet from personally-owned devices, and current and emerging technologies that may be able to mitigate some of those threats. Based on the capabilities of current computer security and voting technologies, the following three issues remain to be significant challenges faced by remote electronic voting systems.

First, remote electronic absentee voting from personally-owned devices face a variety of potential attacks on voters and voters' personal computers. Since the voter's personal computer is outside the control of election officials, it is extremely difficult to protect against software attacks that could violate ballot secrecy or integrity or steal a voter's authentication credentials. These are serious threats that are already commonplace on the Internet today.

Second, remote electronic voter authentication is a difficult problem. Current technology does offer solutions for highly-secure voter authentication methods, but these may be difficult or expensive to deploy. Personally-owned computers may not be able to interface with these methods, such as having the necessary smart card readers for cryptographic authentication using Common Access Cards or Personal Identity Verification cards.

Third, it is not clear that remote electronic absentee voting systems can offer a comparable level of auditability to polling place systems. Because of the difficulty of validating and verifying software on remote electronic voting system servers and personal computers, ensuring remote electronic voting systems are auditable largely remains a challenging problem, with no current or proposed technologies offering a viable solution.

Many of the current and emerging technologies identified in this report are areas with active research and development. Pilot projects should be encouraged, including those involving the use of voting-specific cryptographic protocols, such as the Helios voting system [23]. Emerging trends and developments in these areas should continue to be studied and monitored.

References

[1]   Uniformed and Overseas Citizens Absentee Voting Act (UOCAVA), (as modified by the National Defense Authorization Act for FY 2005). http://www.fvap.gov/resources/media/uocavalaw.pdf

[2]   107th U.S. Congress (October 29, 2002). "Help America Vote Act of 2002 (Pub.L. 107-252)." U.S. Government Printing Office.

[3]   National Institute of Standards and Technology Interagency Report: 7551, *A Threat Analysis on UOCAVA Voting Systems*, December 2008.

[4]   Draft National Institute of Standards and Technology Interagency Report 7682, *Information System Security Best Practices for UOCAVA Supporting Systems*, April 2010.

[5]   Draft National Institute of Standards and Technology Interagency Report 7711, *Security Best Practices for the Electronic Transmission fo Election Materials*, June 2010.

[6]   U.S. Election Assistance Commission (2010, August 10). UOCAVA Pilot Program Testing Requirements, August 10, 2010. Accessed February 16, 2011 at http://www.eac.gov/testing_and_certification/eacs_work_with_military_and_overseas_voting.aspx

[7]   EAC (2010, April 26). Report to Congress on EAC's Efforts to Establish Guidelines for Remote Electronic Absentee Voting Systems. Accessed May 10, 2010 at http://www.eac.gov/program-areas/voting-systems/docs/04-26-10-move-act-report-to-congress-final-congress/

[8]   M. Volkamer and R. Vogt. Basic set of security requirements for Online Voting Products. Common Criteria Protection Profile BSI-CC-PP-0037, Bundesamt für Sicherheit in der Informationstechnik, Bonn, April 2008.

[9]   Council of Europe. Legal, Operational, and Technical Standards for E-Voting. Recommendation Rec(2004)11, September 2004.

[10]   Federal Voting Assistance Program. *Secure Electronic Registration and Voting Experiment. Threat Risk Assessment- Phase 3*. March 23, 2004.

[11] McConnell, Steven (2004), Code Complete (Second Edition), Microsoft Press.

[12] Georgia Tech Information Security Center (2008). *Emerging Cyber Threats Report for 2009.* Accessed May 15, 2010 at http://www.gtisc.gatech.edu/pdf/CyberThreatsReport2009.pdf

[13] US-CERT (2008, May 16). *Technical Cyber Security Alert TA08-137A: Debian/Ubuntu OpenSSL Random Number Generator Vulnerability.* Accessed May 15, 2010 at http://www.us-cert.gov/cas/techalerts/TA08-137A.html

[14] Symantec (2010, April). *Symantec Global Internet Security Threat Report: Tends for 2009.* Accessed May 15, 2010 at http://eval.symantec.com/mktginfo/enterprise/white_papers/b-whitepaper_internet_security_threat_report_xv_04-2010.en-us.pdf

[15] Dierks, T. and Rescorla, E., *The TLS Protocol Version 1.2*, Internet Engineering Task Force, Request for Comment 5246, August 2008, http://tools.ietf.org/html/rfc5246

[16] Atsushi Fujioka, Tatsuaki Okamoto, and Kazui Ohta. A Practical Secret Voting Scheme for Large Scale Elections. In Jennifer Seberry and Yuliang Zheng, editors, *Advances in Cryptology - AUSCRYPT '92*, volume 718 of *Lecture Notes in Computer Science*, pages 244--251, Berlin, 1993. Springer-Verlag.

[17] Rene Peralta. Issues, non-issues and cryptographic tools for Internet-based voting. In *Secure Electronic Voting* (Boston, 2003), Dimitris A. Gritzalis, editor. Kluwer Academic Publishers, pp. 153-164.

[18] Lorrie Faith Cranor and Ron K. Cytron, Sensus: A Security-Conscious Electronic Polling System for the Internet. *Proceedings of the Hawai`i International Conference on System Sciences*, January 7-10, 1997, Wailea, Hawaii, USA.

[19] J. Benaloh and D. Tuinstra. Receipt-Free Secret-Ballot Elections. *Proceedings of the 26th ACM Symposium on Theory of Computing.* Montreal, PQ. May 1994. (New York, USA: ACM 1994), pp. 544—553.

[20] Ronald Cramer, Rosario Gennaro, and Berry Schoenmakers. *A secure and optimally efficient multi-authority election scheme.* European Transactions on Telecommunications, 8:481-489, 1997.

[21]   Premiere Election Solutions (2008, August 19). *Product Advisory Notice*. Accessed May 15, 2010 at http://www.sos.state.oh.us/sos/upload/news/20081001c.pdf

[22]   Fink, R.A.; Sherman, A.T.; Carback, R.; , "TPM Meets DRE: Reducing the Trust Base for Electronic Voting Using Trusted Platform Modules," *Information Forensics and Security, IEEE Transactions on* , vol.4, no.4, pp.628-637, Dec. 2009.

[23]   Ben Adida, Olivier de Marneffe, Olivier Pereira, and Jean-Jacques Quisquater. Electing a University President Using Open-Audit Voting: Analysis of Real-World Use of Helios, In D. Jefferson, J.L. Hall, T. Moran, editor(s), *Electronic Voting Technology Workshop/Workshop on Trustworthy Elections*, Usenix, August 2009.

[24]   Nathanael Paul, Andrew S. Tanenbaum, "The Design of a Trustworthy Voting System," Computer Security Applications Conference, Annual, pp. 507-517, 2009 Annual Computer Security Applications Conference, 2009.

[25]   Common Criteria for Information Security Evaluation. Part 3: Security assurance components. Version 3.1, Rev. 3, July 2009.

[26]   Patrick Peterson, Henry Stern. "Botnets Gone Wild! Captured, Observed, Unraveled, Exterminated." Presented at RSA 2010, San Francisco, CA, March 1-5, 2010.

[27]   Testimony of Bob Carey, Director of FVAP. (2010) EAC Public Meeting, Dec. 3 2009. Accessed April 5, 2010 at http://www.eac.gov/public_meeting_12032010/

[28]   United States Postal Service (2007). *2007 Comprehensive Statement*. Accessed March 17, 2010 at http://www.usps.com/strategicplanning/cs07/chpt5_001.htm

[29]   Alvarez, R. Michael (2005, October 5). "Precinct Voting Denial of Service", *NIST Threats to Voting Systems Workshop*. Accessed March 17, 2010 at http://vote.nist.gov/threats/papers/precinct_dos.pdf

[30]   Davis, Joshua (2007, August 21). "Hackers Take Down the Most Wired Country in Europe" *Wired Magazine*. Accessed March 5, 2010 at http://www.wired.com/politics/security/magazine/15-09/ff_estonia

[31] Markoff, John (2008, August 13). "Before the Gunfire, Cyberattacks" *The New York Times*. Accessed March 5, 2010 at http://www.nytimes.com/2008/08/13/technology/13cyber.html

[32] Vixie, Paul, Sneeringer, Gerry, and Mark Schleifer (2002, November 24). Events of 21-Oct-2002." Accessed March 5, 2010 at http://d.root-servers.org/october21.txt

[33] Internet Corporation for Assigned Names and Numbers. "Factsheet-Root server attack on 6 February 2007." Accessed March 5, 2010 at http://www.icann.org/announcements/factsheet-dns-attack-08mar07.pdf

[34] Worthham, Jenna, and Andrew E. Kramer (2009, August 7) "Professor Main Target of Assault on Twitter" *The New York Times*. Accessed March 5, 2010 at http://www.nytimes.com/2009/08/08/technology/internet/08twitter.html

[35] D. J. Bernstein and Eric Schenk (1996). *SYN Cookies*. 1996. Accessed May 15, 2010 at http://cr.yp.to/syncookies.html

[36] Mell, Peter and Tim Grance (2009, October 7), *The NIST Definition of Cloud Computing*. Accessed March 2, 2010 at http://csrc.nist.gov/groups/SNS/cloud-computing/cloud-def-v15.doc

[37] U.S. Election Assistance Commission (2006, December). *Election Crimes: An Initial Review and Recommendations for Future Study*. Accessed June 15, 2010 at http://www.eac.gov/assets/1/workflow_staging/Page/57.PDF

[38] Gartner (2009, April 14). *Gartner Says Number of Phishing Attacks on U.S. Consumers Increased 40 Percent in 2008*. Accessed March 5, 2010 at http://www.gartner.com/it/page.jsp?id=936913

[39] Cormac Herley and Dinei Florencio, A Profitless Endeavor: Phishing as Tragedy of the Commons, in *Proc. New Security Paradigms Workshop*, Association for Computing Machinery, Inc., September 2008.

[40] Anti-Phishing Working Group (2009). *Phishing Activity Trends Report, 4th Quarter 2009*. Accessed March 5, 2010 at http://www.antiphishing.org/reports/apwg_report_Q4_2009.pdf

[41] Office of Management and Budget (2006, June 23). *OMB Memo M06-16.* Accessed March 5, 2010 at
      http://www.whitehouse.gov/OMB/memoranda/fy2006/m06-16.pdf

[42] McAfee Labs (2009). *2010 Threat Predictions.* Accessed April 13, 2010 at
      http://www.mcafee.com/us/local_content/white_papers/7985rpt_labs_threat_predict_1209_v2.pdf

[43] Department of Defense. *Common Access Card.* Accessed March 5, 2010 at http://www.cac.mil/

[44] National Institute of Standards and Technology (2009). *About Personal Identity Verification (PIV) of Federal Employees and Contractors.* Accessed March 5, 2010 at http://csrc.nist.gov/groups/SNS/piv/

[45] Estonian National Electoral Committee. *Internet voting in Estonia.* Accessed March 5, 2010 at
      http://www.vvk.ee/public/dok/Internet_Voting_in_Estonia.pdf

[46] Mozilla Foundation (2006, November 14). *Firefox 2 Phishing Protection Effectiveness Testing.* Accessed April 5, 2010 at
      http://www.mozilla.org/security/phishing-test.html

[47] S. Egelman, L. Cranor, and J. Hong. You've Been Warned: An Empirical Study on the Effectiveness of Web Browser Phishing Warnings. CHI '08: Proceedings of the SIGCHI conference on Human Factors in Computing Systems. April 2008.

[48] National Institute of Standards and Technology (2008, August). *The 2008 NIST Speaker Recognition Evaluation Results.* Accessed May 5, 2010 at
      http://www.itl.nist.gov/iad/mig/tests/sre/2008/official_results/index.html