Appendix A

California Internet Voting Task Force
Technical Committee Recommendations

1 Scope of the Technical Committee Report

This document is a report from the Technical Committee of the California Internet Voting Task Force. It contains a technical analysis of the communication and security issues inherent in Internet voting, along with recommended privacy and security requirements for any Internet voting systems fielded in California. This report also deals with potential Internet-based voter registration systems and, briefly, with Internet petition-signing systems as well.

We do not describe the design of any particular systems; there is too wide a range of software and infrastructure designs that are potentially acceptable Internet voting solutions and there is every reason to expect that different choices might be made in different counties of the state and in different states. Instead, we recommend *requirements* for such systems, and criteria to be used in their certification, leaving the detailed design to potential vendors.

Because we do not discuss specific designs, we do not include any detailed discussion of costs. They would depend strongly on the goals, design, and scale of the particular system in question. In any case the costs and cost structures in the world of communication and Internet technology are changing so rapidly that an estimate made today might have little relevance by the time such a system is actually procured.

This document is being written January, 2000, and reflects the state of technology as it exists now, or can be reasonably anticipated in the near future. While most of our conclusions are fairly technology-independent, there are inevitably a few concerns and conclusions discussed here that may need revision at some point in the future.

2 General conclusions of the Technical Committee

The Technical Committee has reached a number of general conclusions about Internet-based registration, petition signing, and voting systems. Before detailing all of the reasoning in support of those conclusions, we provide here a quick summary. Each of these conclusions will be expanded upon in later sections.

2.1 Incremental approach to Internet voting

If Internet voting is instituted in California, it should be added in an incremental manner. It should be designed as an *additional* option for voters, not a replacement either for absentee balloting or balloting at the polls; and it should work in the context of the current (paper-based) voter registration system.

Internet voting should, at least initially, remain county-based for greater security and for proper integration with the current registration and voting systems, even though some economies of scale could be realized with a regional- or state-level system.

2.2 Internet voter registration not recommended

The Task Force strongly discourages any consideration of an all-electronic Internet voter registration system. Without online infrastructure for strong verification of the identity, citizenship, age, and residence of the person doing the registering, essentially any all-electronic voter registration system would be vulnerable to *large-scale* and *automated* vote fraud, especially through the possible registration of large numbers of phantom voters.

2.3 Internet petition-signing more difficult to make secure than Internet voting

Besides voting, registered voters in California have the right to formally sign petitions of various kinds, e.g. initiative petitions, recall petitions, etc. Potential systems for Internet-based petition-signing would face essentially all of the same privacy and security issues that arise in Internet voting systems, so most of the recommendations made here regarding security for Internet voting systems apply to any proposed Internet petition-signing system. But because of several structural differences between voting and petition signing that increase the security risks associated with Internet petition signing, we recommend even greater caution be exercised in considering any Internet-based petition signing system.

2.4 Privacy and security issues in voting

Security (including privacy) and reliability are the most important engineering considerations in the design for i-voting systems. Security in this case means (1) voter authentication (verification that the person voting by Internet is a registered voter in the district in which s/he is voting), (2) vote integrity (assuring that an electronic ballot is not forged or modified surreptitiously), (3) vote privacy (assuring that no one can learn how any individual voter voted), (4) vote reliability (assuring that no Internet ballot is lost), (5) non-duplication (assuring that no voter can vote twice), (6) defense against denial of service attacks on vote servers and clients, and (7) defense against malicious code attacks on vote clients.

Reliability means (1) that the entire system, from end to end, operates properly even in the face of most kinds of local (single point) failures; (2) that its performance tends to degrades smoothly, rather than catastrophically, with additional failures; (3) that voters have solid feedback so that they know unambiguously whether their vote was affected by a failure of some kind; (4) the probability of a global system-wide failure is remote; (5) the rarest of all technical failures are those that result in votes being lost after the voter has received feedback that the vote was accepted; and (6) procedures are in place to protect against human failure, either accidental or malicious, that might result in incorrect results of the canvass.

Each of these issues requires specific architectural features (hardware and software) in the design of any system for Internet voting. Most of them are well-understood, with satisfactory technical solutions readily available, which we expand upon in the recommendations below. However some of them require special attention in the case of non-county-controlled (e.g. home or office) voting.

2.5 Internet voting systems should be modeled on the absentee ballot system

The Task Force views Internet voting as being in many ways analogous to (paper) absentee balloting, in that the voter might vote remotely and/or early, and without a personal appearance at the polls. The analogy is even stronger in the case of vote-from-anywhere systems in which the ballot passes through many hands on the way from the voter to the canvass. We therefore recommend modeling some i-voting procedures on established California procedures for absentee ballots, including these requirements:

- A voter must specifically request authorization for i-voting for each election he or she wishes to vote by Internet, authenticated with a hand signature. For systems in which the i-voting machine is run by county officials or county-trained personnel, the request might be made at the voting site immediately prior to voting. For other situations, e.g. home voting (if such a system is ever adopted) the request must be made in advance, *and on paper, not electronically.*
- A voter who has requested i-voting authorization should only be able to vote provisionally at the polls.

- Internet votes must be transmitted in encrypted form and authenticated as coming from a registered voter, much as an absentee ballot must be sealed in an envelope that is signed on the outside.
- Procedures to protect the integrity and privacy of electronic votes during their processing by elections officials should be modeled on those already in the California Elections Code for handling of absentee ballots.

See Section 5.8, Internet voting compared to absentee ballots.

2.6 Two broad classes of i-voting platforms

There are two broad categories of i-voting systems that must be distinguished in any discussion of Internet voting. The difference is based on whether or not the county election agency has full control of the client-side infrastructure and software used for voting:

- *County-controlled systems:* In these systems the actual computers and software used for voting, along with the networks to which they are immediately attached, and the physical environment of voting, are under the control of election officials (or their contractors, etc.) at all times.
- *Vote from anywhere systems:* These are systems intended to support voting from essentially any computer connected to the Internet anywhere in the world, e.g. from home, the workplace, or from colleges, hotels, cybercafés, military installations, handheld appliances, etc. In this case the computers used as voting machines, the software on them, and the networks they are immediately attached to, and the physical surroundings, are under the control of the voter or a third party, but not under the control of election officials.

This distinction is fundamental because with systems that are not county-controlled, the voting environment is difficult to secure against some very important privacy hazards and security attacks that can arise from infection with *malicious code* or use of *remote control software*. Hence, "vote from anywhere" systems must be substantially more complex to achieve the same degree of privacy and security as is achievable with a county-controlled system.

2.7 Four-stage approach to implementing Internet Voting

We recommend a four-stage approach to possible introduction of i-voting in California. Each stage is a technical advance on the previous ones, but provides better service to more voters. These four types of systems are:

(a) *Internet voting at voter's precinct polling place:* Internet-connected computers are deployed at regular precinct polling places alongside traditional voting systems on election day. Voters identify themselves to clerks as usual with the traditional system, and then have their choice of voting methods. Each vote cast on the voting computers is transmitted directly to the county.

(b) *Internet voting at any polling place in the county: Systems of this type are similar to (a), except that* the voter need not show up at his or her own precinct polling place on election day, but may vote at any county precinct polling place equipped for i-voting, or at any other polling place the county might set up at shopping centers, schools, or other places convenient to voters. Non-precinct polling places might be open for early voting for days or weeks in advance of election day, possibly with extended hours. Such sites would still be manned by county personnel, but they would have to have access to the entire voter roll of the county to check registration and prevent duplicate voting, rather than just the roll for one precinct. This might itself be implemented by Internet access to the county's voter registration database.

(c) *Remote Internet voting at county-controlled computers or kiosks*: Systems of this type are similar to (b) except that the polling places should not have to be manned by trained county personnel, but only be responsible lower-level clerks whose job is to safeguard the voting computers from tampering, restart them when necessary, and call for help if needed. A voter would request Internet voting authorization by mail (as with absentee ballots), bring that authorization to the polling place, and then use it to authenticate themselves to the voting computer just before actually voting.

(d) *Remote Internet voting from home, office, or any Internet-connected computer:* These systems permit voting from essentially any Internet-connected PC, anywhere, including home, office, school, hotel, etc.. As with (c), voters would request Internet voting authorization in advance. Later, when it is time to vote, they must first secure the computer against malicious code and remote control software somehow, then connect to the proper county voting site, authenticate themselves, retrieve an image of the proper ballot, and vote.

The first three of these system types are "county-controlled systems", as defined in Section 2.6. We believe that these systems can reasonably be deployed, at least for trial purposes, as soon as they can be built and certified as satisfying not only the current requirements of the California Elections Code, but also the additional requirements we recommend in this document. If the current Elections Code is found to contain language or provisions that prohibit Internet voting, then the legislature will have to act before any trials can occur in which the votes actually count.

The last type of system, (d), is in the category of "vote from anywhere" systems as described in Section 2.6. We do not recommend deploying these systems until a satisfactory solution to the malicious code and remote control software problems is offered.

3 Internet voter registration

Voter registration systems are the basis of election legitimacy in most of the U.S. In most states each county maintains a database of names, addresses, and signatures for all eligible voters in that county who wish to vote. Its purpose is to guarantee that only people eligible by law to vote in a given district can do so, and that no one can vote more than once ("one person, one vote"). Any major compromise of the voter registration system could lead to fraudulent elections.

3.1 The current California voter registration system

To be eligible to vote in a particular district in California a person must be a resident of that district, a U.S. citizen, at least 18 years old, and not in prison or on parole for conviction of a felony. When a person registers to vote, his or her name and residence address are added to the database of eligible voters and he or she is also assigned to a voting precinct and to the appropriate election districts (assembly district, state senate district, congressional district, school district, utility district, etc.). A voter's registration remains valid for all subsequent elections until the county receives information that the voter has moved, or died, or otherwise become ineligible to vote. The voter's handwritten signature is kept on file and is checked against signatures submitted on requests for absentee ballots, on absentee ballot return envelopes, on initiative and other petitions, and, if our recommendations are accepted, on requests for authorization of i-voting.

Today, voter registration in California is based essentially on the honor system. A potential voter simply fills out and mails a voter registration form with his or her name, address, and signature. By signing the form, the voter attests under penalty of perjury to the truth of the name and address provided, *and* to his or her eligibility to vote (citizenship, age, etc.). A potential voter need not appear in person (as one must in order to get an initial driver's license or passport), nor is he or she currently required to present any documentary evidence either of identity or of eligibility to vote. Other than checking that the address listed on the registration form is a real address, and that the post office will deliver to the voter at that address, there is little that a county can do in California to check the legitimacy of a voter registration.

Unfortunately, the current paper-based voter registration system in California carries a potential for at least small-scale vote fraud. Anyone who is willing to fill out, sign, and mail a number of registration forms with distinct false names and real addresses, and who is willing to sign false affidavits, can attempt to register any number of fake voters and subsequently vote multiple times by absentee ballot using those false identities. But the current registration system involves actual paper forms with live signatures, and human inspection of the forms, and so any attempt to commit *massive* fraud successfully by registering a *large* number of ineligible or non-existent voters would be a complex, risky task. Patterns in the false

names or addresses, or the postmarks, or the timing, or the purported signatures, would almost certainly be noticed by local officials, and the fraud would be detected.

A more secure voter registration system would increase the complexity of the registration process, for example by requiring the voter to appear personally before an official, or present documents, or both. This would reduce the voters' convenience, and possibly intimidate some, which together might reduce the number of people who register and vote. The registration process could less intrusively require voters to include additional information such as their driver's license or a portion of the social security number to help improve accuracy. The California Legislature, in enacting the Election Code, has in effect weighed the risk of fraud versus the risk of reduced voter participation and decided that a certain risk of small-scale fraud is worth taking in order to make voter registration a more convenient and less intimidating process for the law-abiding. This committee is not charged with judging the Legislature's decision on these issues and takes no position on the frailties of current paper-based registration system.

3.2 What is Internet voter registration?

There are various systems that might be referred to as "Internet voter registration". Some "print your own registration form" systems use the Internet simply to get a blank registration form to the voter – a service currently provided by the California Secretary of State. Other possible systems might involve registration kiosks of various kinds, and use the Internet to transmit a scanned image of the paper registration form to the county to avoid postal delays and to speed the county's processing of the paper forms. Finally, one can imagine a completely paperless system that would allow voters to register (or re-register) entirely online from a county controlled kiosk or from a home or workplace PC connected to the Internet, without any paper form at all. This is the most ambitious idea, and the most risky. We will discuss these three types of systems in turn.

3.2.1 "Print your own registration form" systems

There are already online services that allow voters to register by bringing an image of the registration form from a server to their PC screens, printing it on their own printers, and then filling it out, signing it, and mailing it, exactly as they would a pre-printed form obtained from the county or state. California already has such a system in place for the federal version of the voter registration form.

One potential problem with such a system is that it is possible that third-party sites might give out registration forms that are not legally correct, for example by not requesting all legally required information, or by failing to inform the voter that a live signature is required. The best solution to this problem is for the state to recommend that third-party sites link to the state site rather than provide their

own versions of the form. That way, when and if the form changes, there will not be a confusion of sites offering out-of-date versions.

"Print your own form" systems amount to allowing a facsimile of the official pre-printed registration form to be used instead of the real thing. As long as the paper registration system remains on the honor system in California, and does not require personal appearance or documentation of eligibility, "print your own form" systems present no difficult security problems. This task force recommends that they be encouraged.

3.2.2 Paper-based registration kiosks

Another type of Internet voter registration system would be an online registration kiosk provided by the county in convenient public places. A voter would fill out the same paper registration form as usual. But immediately, at the kiosk, some of the information would be keyboarded onto an electronic form, and the signature from the paper form would be scanned. The electronic form, along with the scanned image of the signature, would be transmitted to the county by Internet and immediately added to the county's voter database. The original paper form would be transported to the county later so that the paper form with live signature can be on file along with all other registrations.

A kiosk system might be valuable in states where voters are permitted to register up to a time very close to the election, or even on the same day as the election, because it allows the county voter rolls to be updated instantly, without staff labor, and from a kiosk site convenient to the voters.

There are a few potential problems that must be handled. First, the paper forms must still be used and must be reliably transmitted to the county, or the county could be faced with a registration that has no live signature to back it up. Since a scanned image of a signature alone is not a strong enough basis for future identity checks, the registration should not be considered complete until the county has the original signed form in hand. Until such time, the voter should only be permitted to vote provisionally in any intervening election, and the provisional vote should not count in the final tally unless a signed registration form arrives.

Unattended registration kiosks are conceivable. The voter could fill out and sign a paper registration form as usual, and then feed it into a roll-type scanner (as opposed to a flatbed) attached to an Internet-connected computer in such a way that the form is retained after scanning in a sealed box for later retrieval by county personnel. However, paper-handling machines must be treated gingerly, and have a tendency to jam, or feed diagonally; so we believe an attended kiosk will be much more reliable, and certainly much less subject to tampering, vandalism, prank registrations, and user errors such as scanning the back of the form instead of the front.

In theory, potential voters with scanners attached to their own home PCs could simulate a kiosk and do all of the steps of kiosk registration themselves, including transmitting the scanned image of the signed and completed form to the county registration servers, and mailing the original. However, there would have to be standards for the scanning parameters (image format, resolution, color depth) which many users would get wrong; and there would have to be defenses against attacks on the registration servers, whose IP addresses would have to be public. The benefit in convenience to tech-savvy voters with scanners does not seem to outweigh the costs, so we recommend against home simulation of a registration kiosk at this time.

Kiosk-based voter registration systems as described here retain the live signature feature of the current paper system in California, and are essentially automation aids to it. There are no insurmountable security problems with them, so this task force sees no reason why the state should not permit certification and deployment of human-attended Internet registration kiosks.

3.2.3 Security problems in paperless Internet voter registration system

An all-electronic Internet registration system, i.e. one in which a prospective voter can register himself or herself remotely from any Internet-connected PC, without the use of paper forms, seems like an attractive prospect—one that might simplify voter registration and lower its cost. But it is the judgement of this task force that, at the present time, such a system would also be an invitation to automated, large-scale vote fraud, and hence *we recommend that no system for all-electronic voter registration be certified*. This conclusion could be revisited if some kind of national identification infrastructure were created; but an infrastructure that could at least verify the identity of potential voters and some of the criteria for eligibility to vote is not likely to exist in the U.S. in the foreseeable future.

The following discussion explains the reasoning behind this recommendation. A fully satisfactory Internet voter registration system should verify the following:

a) *identification:* make sure that all registrations are associated with a real, living person, not a fake identity or the identity of a dead person;

b) *eligibility:* make sure that everyone who registers to vote is legally eligible to do so;

c) *non-duplication:* make sure that no one is registered more than once, either under multiple names or in multiple districts;

If even the first of these could be accomplished satisfactorily in an all-electronic system, one might judge the idea worthy of more study. Unfortunately, current technology has no way to accomplish any of these goals well. We discuss them in turn.

Identification: First we should note that current paper-based voter registration systems do a poor job of verifying that the registrant is a real person. This is especially true in California, where one has only to be willing to sign a false affidavit and mail it in order to register a fraudulent voter. One might argue that an Internet registration system with the same limitations as the paper system would at least be consistent with current practice, which is time-tested and reflects tradeoffs between security and convenience that the legislature has deemed appropriate. However, there is a crucial difference: with a paperless Internet registration system, *the possibility of registering fraudulent or ineligible voters can be automated*, and electronic registrations, almost by definition, *will not receive the same human scrutiny* as in a paper system. Anyone with a database of real California addresses, which can be purchased at many software stores, could invent fake names for any number of those addresses, register them to vote from a home PC, and later vote any number of times using those fake identities. Furthermore, he or she could do so remotely, for example from a foreign country, and make it appear that the requests came from many different places, all the while leaving no physical evidence, and perhaps being subject to little or no human scrutiny of the registrations, which would be recorded automatically.

The danger of automated, large-scale vote fraud through fraudulent Internet registrations, possibly committed by persons outside the U.S., is so severe that we believe no system should be certified that does not have strong means of identifying the registrant. Risks that may be quite reasonable with a paper system can become completely unreasonable in an automated system.

But there is today no widely-available, standard way to verify a person's identity over the Internet. There are several general techniques that might be considered, but all have serious limitations:

- *Reference to national identification systems:* One might require someone registering via Internet to include a reference to some other trusted database of certified identity numbers, e.g. birth or naturalization certificate number, or passport number. In business situations it is common to ask for social security number or driver's license numbers as a surrogate for identification. But each of these numbers has its limits as a means of identification, with varying standards for their issuance, and none of them is universal, nor available online to counties for this purpose.

  There simply is no national ID system that can be used as a basis for assuring that false identities are not registered to vote via an Internet registration system. Birth certificates are issued by counties, and generally are not online; in any case they may be difficult or impossible to reliably connect to a prospective registrant as they often contain no biometric information at all, or only baby handprints or footprints.

  Passport and naturalization certificates are issued by the federal government, and are also not online—at least they are not available to counties for voter registration purposes.

Even if there were a universal ID number that one could reference, and even if it could be somehow "checked" online during the Internet registration process, merely asking for such a number is not enough since that would still allow the person registering to report someone else's ID number, or that of a person who has died. A stronger mechanism, one that is actually linked to the person who is at the computer registering, would be required.

- *Digital signatures:* Another approach to identifying people through the Internet is via digital signatures. Citizens would create public-private key pairs and register the public keys with a certification authority. They could then participate in various cryptographic protocols, and could, for example, digitally sign their requests for registration via the Internet.

  However, while a digital signature on a registration request proves that the request came from a holder of the private key, it does not prove that the key has been kept properly private, i.e. that it has not been "shared" with others, or stolen. More importantly, it does not prove that that person has only one such key, possibly issued by different certification authorities. A person with multiple keys might freely register multiple times. And while a certification authority might have a policy of trying to issue at most one key per person, in enforcing that policy it would face the same overall problem we are discussing: how does one verify a person's identity in the U.S., and hence ensure that a person does not create multiple "certified" digital identities.

  A recent legislative proposal by Secretary of State Jones would allow Californians to register a public key with the Department of Motor Vehicles after providing proof of identity. The corresponding digital certificate issued by the DMV could then be used as proof of identification for numerous government transactions, possibly including voter registration.

- *County-maintained biometric database:* The strongest approach would be for the county to create (or subscribe to) a database of identification information, requiring potential Internet registrants to submit some biometric that is repeatable, unalterable, and distinctive enough to prevent multiple registrations, e.g. both thumb prints, or a DNA sample. A handwritten signature is not good enough for this purpose because it can be willfully altered: anyone can produce, and then reproduce, numerous different signatures.

  Unfortunately, such a biometric-based system would not prevent both Internet and paper registration by the same voter, because biometric identification within the traditional registration process might be judged contrary to the National Voter Registration Act of 1993 ("Motor Voter"). And, although some personal computers today are being sold with fingerprint readers, and those devices are likely to become more common, there are still no open standards for fingerprint identification. In any case, many Americans are opposed to allowing government agencies to create additional biometric databases beyond those already maintained. They are concerned that information in other databases

could be combined with that in biometric databases to facilitate tracking their behavior or invasion of privacy. Hence, use of biometric methods for identifying voters must be considered currently infeasible on political/privacy grounds.

Eligibility: Even assuming that we could verify the identity of potential voters, an Internet voter registration system should also verify their eligibility, i.e. determine citizenship, age, legal residence, and that the person is still alive. But just as there is no infrastructure for verification of identity, there also isn't any for verification of eligibility, nor is there likely to be any time soon.

Once again, we should note that the current registration system in California does not require any proof of eligibility to vote other than the voter's affidavit under penalty of perjury (and in fact makes it illegal to require such proof); hence one might argue that the standard of proof of eligibility would at least not be lowered if an Internet registration system also required only an affidavit. However, the possibility that, from a single PC anywhere on the Internet, fraudulent registration could be *automated*, is a new danger not present in current registration systems. Such illegal registrations might very well not be caught. In particular, any real people who are ineligible but who are fraudulently registered by someone else might never know it because, knowing themselves to be ineligible, they might never even try to register.

Non-duplication: It is easy to detect when a person registers more than once using the same identity in the same county, and to either ignore it, or treat it as a re-registration. But to detect if a person is registered to vote in more than one county or state requires cooperation among the 58 California counties, or the 3000 counties in the U.S. As before, the current paper based system is open to this kind of fraud at a small scale; but committing it on a large scale would be a tedious process, probably involving the efforts of many people to fill out enough registration forms needed to succeed. With Internet registration, however, the fraudulent registration process could be automated by a single person, from anywhere in the world, leaving no physical evidence.

California encourages, but does not require, registrants to write their driver's license number on the registration form. That feature helps a great deal to control benign duplication; but it is limited by the fact that it is not required, and that the driver's license system itself does not cover all voters and has its own security holes. In general, strong prevention of fraudulent multiple registrations is only feasible if there is a strong voter identification system.

As if these arguments were not strong enough, there is also the danger that the voter registration process might be interfered with by malicious code infecting the computer used for paperless registration. We discuss these issues at length later under the subject of Internet voting; but all of the potential problems that malicious code can present for Internet voting apply to paperless Internet voter registration as well.

Because under current conditions a paperless Internet voter registration system is so fraught with potential for automated fraud, and because there is no expectation that there will be any movement toward online infrastructure for strong identity verification in the foreseeable future, this task force recommends against adoption of any such system at the present time.

4 Internet Petition Signing

Internet petition signing refers to any system in which voters "sign" official petitions, e.g. initiative, referendum or recall petitions, entirely electronically, with the "signature" and associated information transmitted by Internet to the proper agency, either directly or combined with other signatures. Only registered voters are permitted in California to sign petitions.

The Internet Voting Task Force did not consider Internet petition signing at any great length. Hence, in this report we will confine ourselves to comparing it in principle to Internet voting.

First, we should note that many of the security considerations in the design of Internet voting systems apply with little change to Internet petition signing systems as well--in particular, the fundamental distinction between systems in which the entire end-to-end voting infrastructure is controlled by the county vs. systems in which the voting platform is a home-, office-, or school PC. Systems that would allow online petition signing from a home or office PC are vulnerable to malicious code or remote control attacks on the PC that might prevent the signing of a petition, or spy on the process, or permit additional petitions to be signed that the voter did not intend to sign, all without detection. Hence, for the same reasons that we do not recommend Internet voting from machines not controlled by election officials, we cannot recommend similar systems for petition-signing until such time as there is a practical solution to the general malicious code problem and the development of a system to electronically verify identity.

While there are similarities between voting and petition signing, it is important to note that the two are not identical and they have somewhat different cost and security properties:

- Petition-signing is a year-round activity, whereas voting occurs during a limited time window. Hence, servers and other infrastructure needed to support petition signing would need to be running year-round, instead of just during a time window before election day. This may dramatically increase the total cost of managing the system.
- While it is reasonable to expect voters, for security reasons, to submit a signed request for Internet voting authorization each time before they vote (similar to a request for an absentee ballot), it is not reasonable to expect voters to submit a such request each time they wish to sign a petition. As a result,

voters who wish to sign petitions electronically would likely have to be issued authorization (means of authentication) that are open-ended in time. The longer such authorizations are valid, the more likely it is that some of them will be compromised, or sold, reducing the integrity of the petition-signing system over time.

- Voters can sign any number of petitions in an election cycle. Hence, a compromised authorization to sign petitions would be usable for signing any number of petitions, magnifying the damage to the system's integrity.

5 Internet Voting

Today, registered voters in California cast ballots in public elections either by going to the polls in person on election day, or else by requesting in advance an absentee ballot, filling it out, and sending it back to the county, usually by mail. *Internet voting* would allow voters a third option: to vote electronically, with their ballots transmitted securely over the Internet.

5.1 What is Internet voting?

Internet voting (i-voting) refers to any method of voting in a public election in which the voter's ballot is retrieved via the Internet from a county's vote server, presented to the voter electronically on a computer screen, marked electronically by the voter, and then transmitted back to the vote server via the Internet. There are several variations of i-voting that should be distinguished in any discussion, because they have markedly different security properties.

It is important to distinguish direct recording equipment (DRE) systems from i-voting systems. With DRE systems voters also make their choices on a computer, but only at the polls, only on election day; and the votes are stored in the machine in the precinct for later retrieval by election officials, rather than being transmitted over the Internet one by one as they are cast. DRE systems are electronic alternatives to the well-known mechanical voting machines still in use in some jurisdictions in the U.S., and do not present the more serious security problems we will be discussing here that pertain to i-voting.

5.2 What is the value of Internet voting?

Internet voting is intended as a service to the electorate, so that voters might vote more conveniently. Some systems permit voting from more convenient sites than the precinct polling places. Some permit early voting, for a period of time before election day. Some permit home voting, workplace voting, and in general, voting from anywhere that there is an Internet-connected computer.

The hope is that with added convenience and flexibility, voter participation in elections may increase. In addition, the latency of voting should be dramatically reduced from several days for the traditional mailed absentee ballot to a few seconds for an Internet ballot, allowing remote voters to wait until much later in the campaign before committing their votes. Finally, we may expect that the speed and accuracy of the election canvass may be increased, since all Internet ballots can be counted within minutes of the closing of Internet voting; furthermore there should be fewer ways to spoil ballots and fewer ways to miscount them than with the current paper-based equipment, all contributing to an improved elections process.

5.3 Comprehensive vs. incremental approaches to Internet voting

There are at least two stances one could take toward i-voting: *comprehensive* and *incremental*. A comprehensive approach would involve rethinking all parts of the elections process from an online perspective, with an eye toward fielding a unified system for online (a) voter registration and district assignment, (b) voter pamphlets and sample ballots, (c) candidate-, initiative-, referendum and recall petition signing, (d) ballot production, (e) voting, (f) canvass, and (g) perhaps even registration as a candidate for office. It might include administering electoral systems at the state level to achieve economies of scale, rather than at the county level, as is traditional. And it might be accompanied by recommendations for other reforms in the electoral process.

An incremental approach, on the other hand, starts with the current electoral system and introduces Internet voting in stages, extending its reach as experience is gained and technology improves. It proposes minimal changes to the California Elections Code, and attempts to minimize the costs for the new infrastructure, new training for officials, and public education that would be required. An incremental approach retains the current county administration of elections, so that i-voting might be adopted at different times and in different forms to suit each county's needs. If early county experiences with i-voting are successful, cost effective, and supported by the public, the early systems can be improved and extended to more comprehensive ones later.

This task force has come down firmly on the side of an incremental approach to i-voting. Because large-scale i-voting in public elections has not been tried as of this writing, and because fair elections, and elections perceived to be fair, are so vital to government, it seems prudent that we adopt a conservative stance, modeling the requirements for any Internet-based voting system as closely as possible on the current systems that both the public and election officials understand and trust. Wherever possible we propose that Internet-based voting processes be analogous to those used with paper ballots, e.g. for preventing most forms of double voting; for dealing with the rare double votes that do happen (usually unintentionally); for keeping records to prepare for election challenges; and for preventing election agency

personnel from violating voter privacy or tampering with votes. Internet voting should be an evolutionary, not a revolutionary change in the voting process.

Of course, there are some issues unique to electronic voting with no analog in current paper-based balloting systems, such as communication failures, potential overloading of voting infrastructure, potential denial of service attacks on voting servers and clients, and potential malicious code attacks on vote clients. We will make detailed recommendations on these issues.

5.4 Strawman architecture for i-voting system

Figure 1 represents a possible general architecture for the infrastructure of an Internet voting system. It is presented for illustrative purposes only, to give us vocabulary for talking about i-voting in the rest of the document; it is not a recommendation or expectation that this architecture be strictly followed.

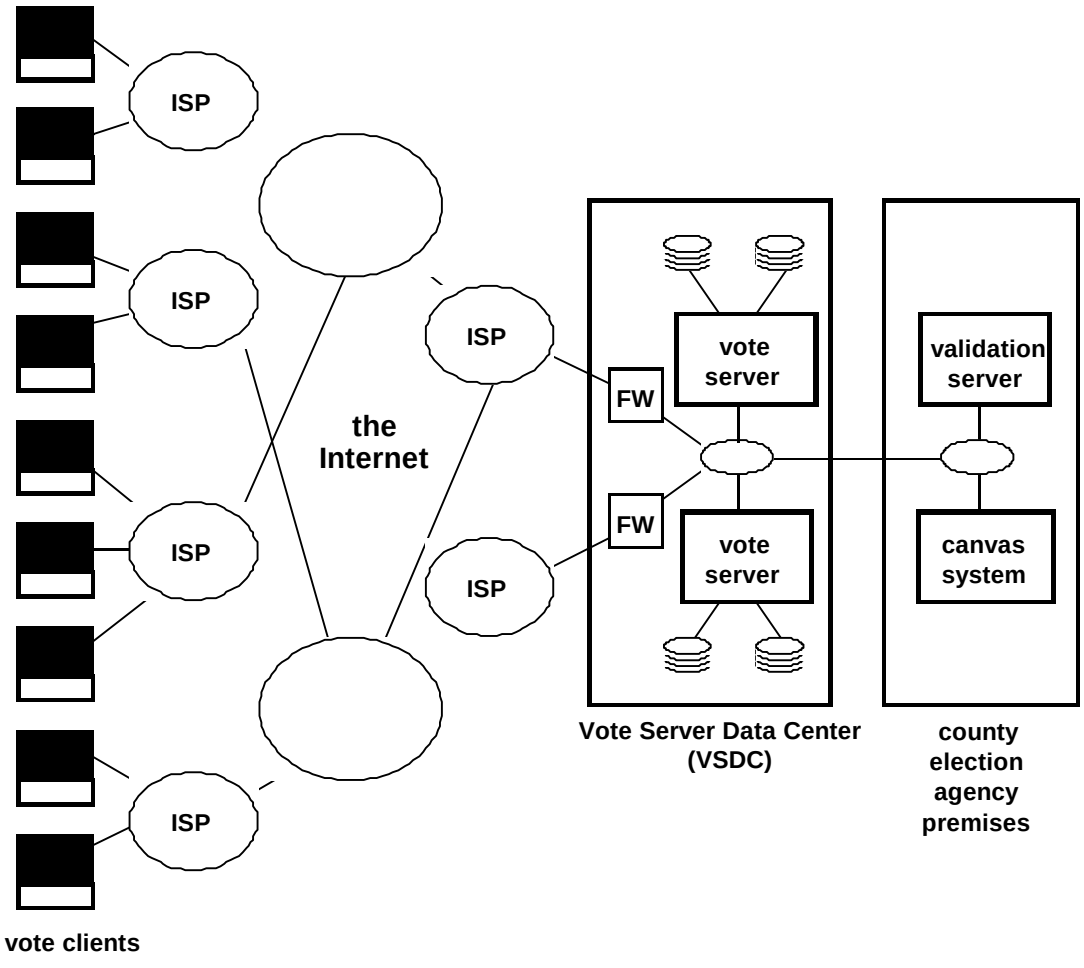

Figure 1: Possible i-voting infrastructure

On the left are vote client machines, i.e. the computers used by voters to cast their ballots. These will generally be small machines (initially PCs of some kind) located in public places such as schools or libraries, or, eventually, in voters' homes or workplaces, etc.

Each client will be connected to an Internet Service Provider (ISP). The ISP's will be connected to other networks that are in turn connected to the ISP's used by the Vote Server Data Center. The complex of ISP's along with the regional and national network service providers they connect to is the Internet. Ballots and related information will travel between the vote clients and the vote servers through the Internet.

We expect (but do not require) that the infrastructure for receiving and counting votes will be divided into two parts, at least logically if not physically. The Vote Server Data Center (VSDC) may be run by the county itself or, perhaps because of the technical skill required to run it, by a vendor under contract with the county. The job of the VSDC is to do the following:

- collect the encrypted electronic ballots from voters submitting them over the Internet;
- store the electronic ballots securely, so that it is essentially impossible to lose any;
- give voters quick feedback that their ballot was accepted;
- transmit the ballots to the county premises for canvassing at some later convenient time

The VSDC, as we envision it, only handles encrypted ballots, and must have no access to any cryptographic keys that could be used to check, read, forge, or modify any ballots. Hence, voter privacy and ballot integrity cannot be compromised at the VSDC without detection. The most vital requirement then remaining is that the VSDC not lose any ballots.

From the VSDC, the ballots, still encrypted, are sent to the county office. This transfer can take place in the background, or just after the close of Internet voting, since high speed is not required.

Canvass of the Internet ballots can be done at the county election offices in a way that is analogous to the handling of paper absentee ballots. Although procedures vary from county to county, in the case of absentee ballots it generally involves checking the signature on the ballot envelope against the signature on file for the voter in the registration records, and checking the database of voters who have already voted. If for some reason a vote has already been recorded for that voter, then the absentee ballot is saved, but not counted; but if not, then a notation is made in the database that he or she has now voted, and the ballot is removed and separated from the envelope. The ballot is put in a pile with other ballots for counting, and the envelope is saved for cross-checking and audit. Once the ballot is separated from the envelope, it is never again possible to match a ballot with the voter who cast it.

In the case of Internet ballots, a similar procedure is necessary to verify that the ballot came from a registered voter from whom no other ballot has been received. The ballot must somehow be tied beyond any reasonable doubt to the voter's registration form, but different i-voting systems will accomplish the linkage differently. It may involve checking the voter's digital signature, or comparing a digitized biometric of some kind to a stored biometric key, etc. Once the ballot's legitimacy has been verified, it should be decrypted and separated computationally from the voter's identity so that they cannot be put back.

Once the ballots are separated from the voter identification information, they are ready for counting. Except that it is accomplished by software, this process is little different from counting of other types of ballots.

5.5 Classification of i-voting systems

This task force has identified four distinct types of Internet voting systems that we believe will work in California. They can be placed in a sequence of increasing complexity leading from relatively simple systems providing modest new services to the electorate with few security concerns, all the way to very sophisticated systems providing unprecedented new convenience to voters, but with more complex security issues to be overcome. These four types of systems are:

(a) Internet voting at voter's precinct polling place;
(b) Internet voting at any polling place in the county;
(c) Remote Internet voting at county-controlled computers or kiosks
(d) Remote Internet voting from home, office, or any Internet-connected computer

While the space of i-voting systems can be sliced in other ways, this classification has the virtue of suggesting a long-term implementation strategy as well: the simpler systems can be implemented first, and the more complex ones can later be built upon the foundations of the earlier, simpler ones when the technology is ready.

In the next four sections we describe these types of i-voting systems in a little more detail.

5.5.1 (a) Internet voting at voter's precinct polling place

The simplest i-voting system is basically a computer set up at precinct polls on election day as an alternative voting device to whatever system is traditionally employed by the county. Voters would enter the polls on election day and identify themselves as usual to poll workers; then they would choose to vote using either the traditional system is employed in the county, or one of the Internet voting terminals. (Eventually some counties may eliminate the traditional voting methods, but that would be very unwise in the first few election cycles because of the possibility of problems with or failures of the Internet systems.)

Such a system provides only modest service to voters, because they have to come to the precinct polls to take advantage of it. It's main benefit is to speed the vote canvass, since the votes are transmitted directly to the county instead of being held in the machine for transmission after the close of the polls. It will likely also have great value as a first step in the construction of more complex systems.

5.5.2 (b) Internet voting at any polling place in the county

In this type of system the county sets up voting computers at places that might be convenient for voters around the region such as shopping centers, schools, town centers, and locations near large employers. County A might even be locate polling places in a neighboring County B if that would be convenient for voters registered in County A. These new sites would be in addition to the traditional precinct polls. Like precinct polls the new sites would be manned by election officials or poll workers, but unlike precinct polls, any voter in the county could vote at any of these sites. Furthermore, the sites might be available for voting in advance of election day as well as on election day, perhaps for several weeks, i.e. as long as the absentee balloting window is open.

Voters would identify themselves to poll workers at these sites exactly as they would at a precinct poll site, but the poll workers would have their own computers with Internet access to the county database of registered voters so they could verify eligibility, determine which ballot style the voter should get, and record that the voter has voted. The poll worker would then give the voter a code of some kind to take to the i-voting computer, both to authenticate the voter to the i-voting computer and to retrieve the proper ballot type.

5.5.3 (c) Remote Internet voting at county-controlled computers or kiosks

This type of system is quite similar to (b) above, except that the voting sites need not be manned by official poll workers. Instead, the i-voting machines at the new polling places, perhaps enclosed in kiosks, would be tended by people with lower-level skills whose responsibility would be only to prevent tampering with the machines, prevent electioneering, prevent voter coercion, and to call for help if any problem develops.

For these systems to be secure, voters would have to have previously requested Internet voting authorization from the county, on a paper form with a live signature, much as voters may now request an absentee ballot. The county would return to the voter a code to be used at the time of voting, both to authenticate the voter and to enable retrieval of the proper ballot type. Presumably this code would be similar to that given to the voter by a poll worker in systems of type (b). Then, in order to vote, voters would simply walk up to an i-voting machine, authenticate themselves using the code provided by the county (without talking to any poll worker), make their choices, and transmit the ballot.

After voters get used to them, systems of this type should be lower in cost in the long run than those of type (b), because they do not require fully-trained poll workers to supervise them. They should therefore be of greater service to voters because presumably more voting sites could be fielded.

5.5.4 (d) Remote Internet voting from home, office, or any Internet-connected computer

Systems of this type allow voters to vote from essentially any Internet-connected computer (with appropriate software) anywhere, including from PCs at the voter's home, workplace, school or college, hotel, or even possibly from a voter's handheld Internet appliance, etc. As with systems of type (c), voters will be required to request authorization for this type of voting in advance, so they can be given credentials (of some kind) by the county for use at the time of voting. In some systems it might be necessary for voters to be issued voting software as well and may also include provisions for the voters to provide the county with a personal identification number (P.I.N.) to be used for voting purposes.

These systems would provide by far the greatest convenience to voters, who could, in effect, vote any time, anywhere. But these systems also involve much more difficult security problems since the election agencies will not have full end-to-end control of the infrastructure for voting.

5.6 County-controlled iVoting computers

For county-controlled i-voting computers, used in systems (a), (b), and (c) above, the most difficult security issues, malicious code and remote control/monitoring software, can be effectively avoided by running a "clean" copy of a stripped-down, minimal operating system and voting application. The software should come directly from a certified source on read-only media, and no software modules or functionality should be included beyond the minimum necessary for i-voting. No remote control or monitoring software should be loaded, nor any software for email, chat, audio (except perhaps in service to blind or illiterate voters), video, file transfer, printing, general web browsing, or other network services extraneous to voting. There should be no software for sharing files or devices over the network, and except for booting the operating system and launching the voting application, it should be possible to do without a file system at all! Unnecessary software that cannot be practically removed for some reason should be turned off or otherwise disabled. Since many of these features tend to be built into the operating systems or browsers of today, it may take some effort, and possibly the cooperation of software vendors, to procure a software base suitably stripped-down for voting. The details should be examined carefully at the time a system is presented for certification.

The most serious remaining issue is tampering. County-controlled machines might in some situations be in service for up to several weeks prior to election day, might be physically handled by hundreds of voters per

day, and might be unused during nights or weekends. A vendor of voting systems intended for use in a public place should provide the specific software configuration intended for that environment, and specific security and maintenance procedures to make sure the machines remain secure. Furthermore, the systems themselves should always be monitored by someone whose job it is to prevent tampering. Other anti-tampering precautions should be considered as well, such as:

- configuring the software so that it requires a password to boot;
- disabling access to the "desktop" so that under no circumstances can the voter can do anything other than vote from the machine;
- configuring the unit, e.g. with cabinetry, so that the voter has physical access only to the screen (and perhaps to a keyboard and/or pointing device if it is not a touch-screen), leaving all other parts inaccessible, especially devices such as floppy drives, CD drives, and any others from which a tamperer might be able to reboot or install software; and
- configuring the machine so that it has no modem, network Interface, wireless communication devices, etc. other than the one needed to connect to the Internet.

5.6.1 Voting from home, the workplace or other institutional computers

The most serious problem in home environments is the possibility that the home PC might be "infected" with a malicious program designed specifically to interfere with voting. Home PCs are generally not professionally managed, and most home users are either not aware of security hazards that might affect voting, or may not know how to use the security tools available. As a result, their computers are frequently vulnerable to all kinds of malicious code attack. For more discussion of this problem, see Section 6.2, Malicious software.

The only way that home voting can be made safe is to have the voter deliberately secure his or her computer just before voting. There are a number of ways to accomplish this with current technology, but all of them require some inconvenience to the voter and some development complexity on the part of the i-voting vendor. See Section 6.2.2, Internet voting systems designed to thwart malicious software.

In the home setting, there is also some risk of loss of voting privacy, since one person might be able to spy on the voting of another. However, we believe that voters at home computers might be presumed to trust other people in the same household. While people might be able to spy over each other's shoulders during voting, or monitor one computer from another on the same home network during voting, people can also spy on others filling out an absentee ballot, or steal each others' absentee ballots. Voters must take *some* responsibility for guarding the privacy their own vote, and the household seems a reasonable boundary within which to expect them to take that responsibility.

In an institutional setting, where the network and the computers are owned and managed by someone other than the voter, it is usually the case that the computers must have a full complement of operating system and networking software for their primary mission. Although they are often just as vulnerable to malicious code attacks as home machines, a "clean system" approach, with an explicit step of securing the platform before voting, may not work well in a workplace environment because rebooting from a clean operating system would likely make the machine unavailable for its primary business purpose.

In addition, workplace voting introduces a new major concern about vote privacy. Institutional computers are often maintained, managed, and controlled by professional staff, rather than the primary user. They are likely to have remote control or monitoring software in place, which leads to the possibility of one employee surreptitiously monitoring (electronically) another's voting. Vendors who expect their i-voting systems to be used in the workplace must go to some lengths to ensure that voter privacy is not compromised. Furthermore, voters in general should be educated about the fact that computers located in places where the security environment is totally unknown, or not trusted, are probably too risky to be used for i-voting. This would include other people's homes, institutions, cybercafes, etc.

Institutions often have their internal networks separated from the Internet at large by a firewall that strongly restricts the kinds of traffic that can flow in and out. Yet another complication that vendors will have to deal with if they expect people to vote from workplace computers is to design their voting system to be compatible with the firewall configurations routinely in use.

Our discussion so far has tacitly assumed that the voting platform is a PC of some kind (including the Apple Macintosh). But new Internet-capable devices are beginning to appear, e.g. hand held electronic organizers, cell phones, "wearable computers", and perhaps "network computers" (NCs). These devices all have substantially different operating systems, screen sizes, and "browser" software than today's PC platform does. It is not likely that an Internet voting system that works from the PC platform will also work from all of these other platforms, at least without substantial adaptation. One risk in the design of Internet voting systems today is that the era of approximate uniformity in the technology base used for interacting with the Internet that is caused by the near ubiquity of the "Wintel" architecture will some day break down, and there will be no clear choices of platform from which to support voting. Vendors and counties should pay attention to this possibility before investing heavily; it is one of the risks caused by the speed of technical change.

5.7 Steps in Internet voting

Internet voting, as we envision it, proceeds in the following sequence of steps, as viewed from the perspective of a voter. Different i-voting systems that satisfy our overall requirements may vary from this in detail, but will generally resemble the following outline:

Voting preliminaries:

1. **Registration**: The potential voter must register to vote. Except in a few special cases the signature on the request must be a live ink signature, and is the primary authenticator used to verify the right to vote, request an absentee ballot or Internet balloting authorization, or sign a petition.
2. **Request for Internet balloting**: Prior to voting the voter may request Internet balloting, on a form similar to the request for an absentee ballot. The request may be delivered to an election official in person or sent by mail, and must include a live ink signature to match against the voter registration record. Hence, a request cannot be accepted by email. A voter should not be able to request both an absentee ballot and i-voting and then choose later which to use.
3. **Authorization:** The county responds to the request, sending the voter, probably by U.S. mail, information about how to authenticate himself/herself and vote online. The information sent and the procedure to be used by the voter will differ with different Internet balloting systems. The voter is marked as having requested Internet balloting, so that if the voter shows up at the polls to vote, he or she will be given a provisional ballot rather than a standard ballot as a guard against double voting.

Voting:

4. **Securing the voting platform:** If the voter is voting at a county-controlled site, or from a secure special purpose device, then there is nothing to do in this step. But if the voter is voting from his or her own computer, or one belonging to a third party, then some steps may need to be taken to secure the computer against malicious code or against third parties monitoring the voting process. Precisely what must be done depends on the design of the specific i-voting system provided by the vendor, but it may involve rebooting the computer in “safe mode”, or from a special county-provided CD-ROM, or it may involve attaching a special device to the computer, etc.
5. **Authentication and ballot request:** During the time window for i-voting, a registered voter with authorization for Internet balloting can vote by Internet. When the voter wishes to cast an Internet ballot, he visits the Internet balloting web page for the proper county and authenticates himself to that server according to the procedures given in step 3 and requests a ballot in the language of his choice. The precise mechanics will differ from one voting system to another. County-controlled voting computers will likely be configured to do nothing but run the voting application and connect to the

county voting site, whereas at a home or workplace PC one might have to deliberately run a browser or voting application and connect to the voting server before authenticating oneself.

6. **Ballot delivery:** The server will send back to the voter an image of the appropriate ballot for his or her precinct in the language requested.
7. **Voting:** The voter marks the ballot with the keyboard and mouse (or touch-screen, if equipped).
8. **Transmission of ballot:** When the voter is finished making choices, he or she clicks a button to send the ballot (and then confirms it again). The ballot is encrypted and sent to the vote server. All unencrypted record of the ballot is then erased from the voter's computer.
9. **Acceptance and Feedback:** The vote server accepts the vote and sends feedback to the voter acknowledging that the vote has been accepted.

Processing the ballot:

10. **Validation and anonymization:** The vote is validated as being from a legitimate voter who has not yet voted, separated permanently from the identification of the voter, and stored for counting.
11. **Verification:** The voter is finished, but may return later to the county web site to check that his or her vote has not only been accepted (i.e. stored), but also authenticated (i.e. validated as a legitimate vote), and will thus be entered into the canvass (i.e. counted). However, the voter cannot, under any circumstances, retrieve a record of *how* he or she voted, or change his or her vote once the ballot is cast.
12. **Canvass:** The votes are counted
13. **Audit, recount, contest:** The votes, the separated identifications of the voters, along with other information, are retained for later audit or recount, or for evidence in case the election is contested.

5.8 Internet voting compared to absentee ballots

This task force has been consciously guided by experience with absentee balloting in the design of requirements for i-voting. In many ways Internet votes, as we conceive them, can be thought of as the electronic equivalent of paper absentee ballots. Both allow ballots to be cast remotely, in principle from anywhere in the world, and at any time convenient to the voter within a time window in advance of election day. With the current California voter registration process, there are inevitably similar procedures for requesting absentee ballots and i-voting authorization, similar mechanisms for prevention or detection of double voting, similar concerns about lost ballots or lost authorizations for i-voting, and analogous mechanisms for protecting ballot secrecy.

But similar as they are, there are some important differences between the two. One is that i-voting systems can give immediate feedback to the voter that his or her ballot was received and accepted; with absentee ballots sent through the mail there is no automatic indication to the voter that it arrived, or arrived on time. There are also ways of spoiling ballots, or over-voting with an absentee ballot, that have no analog with electronic ballots. But the most important difference is that there are security issues arising in i-voting that have no analog in the absentee ballot system. Much of this document will be devoted to discussion of these security issues.

5.9 Elections conducted at the county level

In the U.S. almost all public elections, whether municipal, county, state, federal, or other (e.g. school or utility districts), and whether primary, general, or special, are conducted by county governments. On major election days there are thus 58 parallel elections in California, with the counties reporting the results of state- and federal-level contests to the Secretary of State's office in Sacramento, and the results of other contests to the appropriate officials in those jurisdictions.

Each county, based on its history and needs, makes its own choice of voting systems from among those certified by the Secretary of State. Most counties in California today use a punch card system. A large number of others use one of two mark-sense card systems. In the past, various counties have used mechanical voting machines. And recently several systems for voting at a computer-controlled touch screen and keyboard have been certified for use in California and are now being used by several counties. All counties in California permit absentee ballots as well. Internet voting systems would, from one point of view, be just another voting system.

It is tempting to recommend a system of i-voting to be administered at the state level, since there are substantial communication and computational economies of scale that could theoretically be achieved at that level. But barring major changes in the Election Code, Internet ballot types will have to be assembled and edited in the same way as paper ballot types (with sometimes hundreds of distinct types in up to six languages in one county). And Internet votes will still have to be aggregated with paper votes in contests at all jurisdictional levels. Currently the counties are set up to handle these complications, so it would greatly increase the logistical complexity of elections if i-voting were conducted at any level other than counties when the rest of the system is still county-based.

There is a strong security advantage as well to conducting Internet voting at the county level. If a uniform statewide system of i-voting were adopted and widely used, then certain security attacks, such as malicious

code attacks against voters' computers, or denial-of-service attacks against vote servers, could be much more effective, possibly swinging the results of statewide elections or electoral votes in a presidential election. Such a circumstance may be much more tempting to someone with a motive to interfere with an election. However, if i-voting is adopted at the county level, and different counties adopt different systems, or variations on the same system, and some counties do not adopt it at all, then a potential attacker has a much more difficult problem. Any single attack scheme is likely to work only in one county, or a few counties with nearly identical systems, with a corresponding reduction in payoff to the attacker. County-level attacks may not be worth the risk of jail to an attacker, whereas a state election conceivably might. Diversity in i-voting systems around a state, like genetic diversity in a biological system, tends to protect against large scale attacks against the system as a whole.

We therefore assume that any i-voting systems will also be administered at the county level. Each county should have the authority to choose, based on local circumstances, from among the set of i-voting systems certified by the Secretary of State. Some counties will adopt i-voting systems earlier than others; some may reject i-voting entirely; and conceivably some might adopt more than one i-voting system for any of a number of reasons, e.g. to give voters a choice, or because a more streamlined system is appropriate for some local or special elections.

6 Security in i-voting

The current paper ballot systems set a security standard that we adopt as the baseline for i-voting. They represent certain tradeoffs between voter convenience and protection against fraud that the Legislature and Congress, have deemed appropriate; hence we take it as a guiding for the design principle. We require that elections with i-voting be at least as secure as those without; however, we view our charter as not to make broad recommendations for election security reform, but to offer means to integrate i-voting as smoothly as possible into the current systems.

In any engineered system there are design tradeoffs that reflect necessary compromises between conflicting goals. In i-voting, one key tradeoff is between ease and simplicity of voting on the one hand, and the integrity and privacy of votes on the other. Absentee balloting, for example, is more complicated than voting at the polls, even though it is potentially less secure. The requirement for voters to send a *new request* for an absentee ballot for each election, and do so with a live signature, and then sign the ballot envelope when mailing it back, are all security procedures that have no analog when voting at the polls, but are the necessary price to be paid for the convenience of remote, early voting afforded by absentee ballots. Likewise, i-voting will have its own security procedures, which will often make voting more complex than other Internet transactions, more complex than voting at the polls, and, when voting from home, school, or

office PCs (as opposed to a voting kiosk), more complex than using a paper absentee ballot. The additional complexity is the inevitable price of security and convenience.

Since i-voting systems are assumed here to augment, rather than replace, voting at the polls and voting with paper absentee ballots, this task force has adopted the criterion that *the overall security of elections must not be reduced by the addition of i-voting as an option.* But in the absence of improvements in security of the current registration and voting systems, a very tight security for Internet voting can do little to increase the overall security of an election. Putting strong locks and guards on one barn door, when there are weak locks and no guards on the other doors, does not increase the overall security of the barn.

As an application of this reasoning, we note that there are some weaknesses in current electoral practice that we do not anticipate will be rectified in I-voting systems. Among them are the potential for vote coercion, or the sale of votes, or potential privacy violations under the current absentee ballot system. Nothing prevents a voter, perhaps under coercion, from allowing another person to watch over his shoulder as he votes and mails the ballot. Nor does anything prevent him or her from pre-signing the ballot envelope, thereby authenticating it, and then selling the envelope and the blank ballot to someone else who then casts the vote (other than the fact that it is illegal). Neither of these problems occurs with voting at the polls. Since these possibilities are already inherent in the current absentee ballot system, we did not adopt the criterion that they must be prevented with i-voting systems.

On the other hand, we did not want to introduce *new* modes of vote coercion or vote sale, or extend their scope or time window. For example, several security problems could be solved or ameliorated if it were possible for Internet voters to contact the county after voting to verify how they voted—a possible feature that is perfectly feasible technically, but has no analog in paper voting systems. However, that would also allow the coercion or sale of votes not just *before* the ballot is mailed, but also for as long *afterward* as the window of verification remains open. We believe that would open the door to widespread abuse, and would reduce the overall security of elections; hence, we recommend instead that there be no way for an Internet voter to verify his or her vote after the fact.

6.1 Security issues specific to i-voting

There are several broad security issues that must be dealt with in any i-voting system that are specific to Internet voting, and may have no analog in conventional voting systems. Here is a short list of them:

- *voter authentication:* determining that a ballot arriving at the vote server really is from the registered voter it purports to be from;
- *ballot privacy:* preserving the secrecy of the ballot—that no unauthorized person can read the ballot, and no one can associate a ballot with the person who cast it;
- *ballot integrity:* guaranteeing that ballots cannot be surreptitiously changed by any software agent or third party;
- *reliable vote transport and storage:* guaranteeing that no ballot is either created or destroyed (lost) anywhere from the vote client to the vote server without detection, and no ballots at all are created or destroyed (lost) at all from the vote servers to the vote canvass computers;
- *prevention of multiple voting:* no more than one ballot may be counted for any one voter;
- *defense against attacks on the client:* guaranteeing that there is no malicious software (Trojan horse, virus, etc.) on the client that can affect the integrity or privacy of the ballot;
- *defense against denial of service attacks on vote servers:* dealing with deliberate attacks intended to control, crash, or overload the vote servers or the networks they are attached to.

The first four of these properties are referred to as “end to end” properties, in that they call for maintaining a security property all along the multi-step path from one end of the communication (the mind of the voter), to the other (storage on the county vote servers or canvassing computers). For example, ballot integrity requires that the contents of a voter’s ballot not be changed by malicious software on the computer he or she votes on, nor by any of the routers, computers, or employees of the several private networks along the Internet path to the vote servers, nor by the vote servers themselves, nor by any employees of the contractor that runs the VSDC, nor in transit to from the VSDC to the county canvass computers.

If the voter is voting from a home PC, the most insecure, uncontrolled part of the end-to-end path is inside the computer used by the voter. Any i-voting protocol will transmit the ballot in encrypted form, which guarantees that it cannot be read by any third party, and that it cannot be modified by a third party without detection. Therefore, the riskiest part of the trip that the ballot takes is *inside the vote client, before it is encrypted.*

6.2 Malicious software

Malicious software is software that is deliberately designed to do harmful things that the user neither wants nor expects, and to either hide the harmful action or perform it so quickly that it cannot be stopped. Also known as *malware* or *vandalware*, it can be introduced on a client machine, and in such a way that the voter is unaware of its presence. Among the things that malicious software can easily do if no preventative measures are taken are (a) change the votes on the electronic ballot without the voter's knowledge, (b) reveal the supposedly secret votes to some outside party, or (c) simply prevent a person from voting, possibly leaving him or her with the impression that the vote was recorded.

Malicious software is usually distributed to home and office computers through a variety of mechanisms known in the security literature as *viruses*, *worms*, *back doors*, *trapdoors*, *logic bombs*, *Trojan horses*, *bacteria*, *rabbits*, or *liveware*. Prof. Eugene Spafford of Purdue University provides an excellent set of definitions and discussions around each of these methods.

6.2.1 Scope of the malicious software problem

Malicious software is probably the most difficult technical problem involved in i-voting. While we will describe the problem is some depth to indicate its seriousness, it is important to keep in mind that *there are solutions*, some of which we will describe in a later section.

Today's PC operating systems are designed as open software systems, so that users routinely change their functionality by adding device drivers, DLLs, extensions, control panels, patches, upgrades, and other code modules acquired from any number of places. Usually such code is added to the operating system as a side-effect of deliberately installing application software or system upgrades, although operating system changes can also be caused by viruses. In any case users are frequently unaware that the operating system has been changed, and certainly have no way of certifying the safety of the changes.

Browsers are even more open and more casually modified through the addition of such code modules as plug-ins, Active-X controls, JavaScript scripts, and Java applets. In many cases programs are downloaded without the user's knowledge as an invisible side-effect of merely visiting a web page, and yet they have full power to modify the software base and behavior of the computer arbitrarily.

This easy extensibility of the operating system and browser are extremely valuable for the general flexibility and adaptability of PC software. It is part of what allows such astonishingly fast evolution of PC technology. But the background danger is that any of these kinds of software extensions can harbor a malicious program, for example a "Trojan Horse", i.e. a program that surreptitiously does something other

than it is advertised to do, usually harmful in some way to the user's files. Since a typical home PC has numerous operating system and browser extensions from a wide variety of places, and since there is not, and cannot be, a general test for whether these extensions carry malicious code, the home PC is an extremely dangerous platform from which to perform transactions that must be secure.

If voting were permitted from PCs with standard web browsers running over a standard operating system with no further security measures, then it would be very easy for a rogue programmer to write a malicious program in the form of an ActiveX control or plug-in or virus, then lure thousands of users to download that code, possibly *unknowingly*, and have that rogue program either *spy on* the user's voting, or even *change* the user's votes without the voter's knowledge, and regardless of any other features of the i-voting protocol.

A special case of the problem arises with computers connected to local area networks (LANs), or connected to the Internet through certain technologies such as cable modem connections in which the last link of the coaxial cable is, in effect, a local area network connecting many households in the neighborhood. Unless the software on a computer is very carefully configured, it is extremely easy for a person on one computer to install software, including malicious code or remote control software, on another computer on the same LAN. In the case of computers connected to certain cable Internet access systems, this would include computers owned by strangers in other nearby households, whose owners are very unlikely to know this is possible.

It is essential that any i-voting system offer some kind of guarantee that it is immune from the sort of malicious code attack that could affect the outcome of an election. It is not sufficient to argue that such an attack is unlikely, or even *very* unlikely. An election would be an extremely tempting target for any motivated person, from a lone hacker to a political partisan to a foreign government. Such an attack would be a political and public relations disaster; or worse, if undetected, compromise the results of the election. We must presume therefore, that if a malicious code attack is possible, it will happen sooner or later. Even before it happens, security experts will surely criticize publicly any election system having such a vulnerability, and the public would likely lose confidence in such a system.

It is important to understand that the problem of malicious code on PC platforms (including Macs and other computers) cannot be fully solved simply by adding more software, because it is a fundamental fact of the theory of computation that there can be no general test to detect whether or not a PC is harboring malicious software. Commercial virus detection software can detect and neutralize *known* viruses and other malicious programs that have already come to the attention of the security experts. But they can do very little about *unknown* malicious programs, such as those that might be quietly lying in wait for a specific event (e.g. voting) and that then take invisible action (e.g. changing a vote).

There are ways around the malicious code problem, but they all require security measures beyond ordinary use of the current PC platform and browser. It may involve a new operating system with a security architecture built in from the ground floor. It may rely on some device, communication, or human process that occurs *outside* the PC, and would therefore be immune to manipulation by a malicious PC program—perhaps telephone communication, or paper communication via the postal service, or a closed, uninfectable security device that plugs into PC via the serial or USB port. Or it may involve some special-purpose appliance, useful only for voting, that is software-closed and communicates with the Internet directly, bypassing PCs altogether. But i-voting mediated *solely* through standard PCs with the standard software available now or in the next couple of years is not recommended.

6.2.2 Internet voting systems designed to thwart malicious software

As indicated, there are ways to design i-voting systems that detect, avoid, or ameliorate the problem of malicious code. Most of them have in common one crucial point: that all cryptographic operations, and all manipulation of unencrypted vote data, take place in a software context that cannot be affected by malicious code.

Here we enumerate some of the possible approaches to the problem of malicious software; this list is not exhaustive, and other approaches might be created and certified.

1. *Clean operating system and voting application:* Prior to voting, the voter's machine could be booted from a CD-ROM (or similar media) containing a "clean" operating system, with no extensions that might harbor malicious code. Combined with sophisticated scans for an infected BIOS (or equivalent on other computers), this step could virtually eliminate the possibility of malicious software during voting. This is presumably the approach that would be used for county-controlled voting machines; but such a CD-ROM could also be distributed for home voting via the postal service in response to a voter's request for i-voting authorization.

   The application program used for browsing, presumably distributed on the same CD-ROM, would also have to be "clean". Current commercial browsers are not suitable for voting because they are particularly vulnerable to malicious software. A special-purpose web browser that does not accept extensions such as plug-ins, applets, controls, or scripts, and that is dedicated solely to voting, would be far more resistant to infection than today's commercial browsers, and its integrity could be conclusively verified with a cryptographic hash or digital signature.

2. *Special security PC hardware:* A special, software-closed security device might be developed to be attached to the voter's computer, e.g. through a USB port. Its purpose would be to display the ballot to the user, accept the voter's choices as input, and perform the cryptographic operations. In effect the

voting is done on the security device, and the PC it is attached to is used only as a conduit to the Internet. Since the device is software-closed, meaning its software cannot be changed, it is not subject to infection by malicious code.

3. *Closed, secure devices:* It is possible that special, software-closed, Internet-capable devices, such as network computers (NCs) or hand-held, wireless descendants of today's cell phone and electronic organizers, may be developed for commerce and may be secure enough for voting as well.

4. *Secure PC operating systems:* Future commercial PC operating systems may be designed for greater security than today's systems. For example, they may be composed of digitally-signed modules, allowing secure applications to exclude, as untrusted, modules of dubious origin (i.e. potentially malicious programs). Such an operating system would enable practical, secure home and workplace voting.

5. *Code sheets:* Voters could be mailed code sheets that map their vote choices to entry codes on their ballot. While voting, the voter uses the code sheet to know what to type in order to vote for a particular candidate. In effect, the voter does the vote encryption, and since any malicious software on the PC would have no access to the code sheet, it would not be able to change a voter's intentions without invalidating the ballot.

6. *Test ballots:* Special test ballots can be sent from vote clients and checked by software at the county. The number, location, timing, and contents of the test ballots should be known by the county, but they should be otherwise indistinguishable from real ballots, so that any malicious code that destroys or changes real ballots will affect the test ballots as well. Analysis of the test ballots will enable any malicious code attacks to be detected, the locations of infected machines to be determined, the approximate time of the attack to be estimated, and the total number of votes affected to be bounded.

   Note that this technique does not *prevent* malicious code attacks; it only *detects* them after the fact. Hence it must be combined with one of the previous preventative techniques. Still, it is a very powerful technique because it can also be used to detect *any* systematic cause of lost ballots, not just malicious code attacks, and because it provides a quantitative measure of the size of any problem it detects.

7. *Obscurity/complexity:* One final approach, while not sufficient for real security, nonetheless raises the cost to potential attackers. Digital ballot formats and voting software may be kept secret prior to the election and possibly randomly changed during the election, or made complex in other ways. In order to successfully carry out an attack and escape detection, malicious software authors must have a great deal of information about the internal format of the ballot and voting software. If these details are not available in advance, and/or if that information is complex, the potential authors of attack software may not have enough time to develop and distribute it during the election window.

6.2.3 Security for i-voting vs. security for electronic commerce

A commonly-asked voting security question is this: If the PC is widely used for secure electronic commerce over the Internet, and people buy everything from books to stocks online, then what is so problematic about online voting? Aren't the authentication, integrity, privacy, and malicious code concerns similar for the consumer and the voter?

The simple answer is "No". Security issues in i-voting are more difficult than for electronic commerce because of one fundamental difference: in electronic commerce, financial transactions are performed online, but there is a separate offline process for checking them and for correcting any errors detected, whereas such is not, and cannot be, the case for voting. Therefore, the fundamental security emphasis in voting must be *up-front prevention of fraud and error, with no reliance on any possibility of after-the-fact correction,* a much more stringent requirement than is generally necessary today for financial transactions.

Online financial transactions today are usually followed later by account statements delivered on paper from the credit card company or merchant. The consumer should, and usually does, check those statements, at least superficially, and can contact the merchant or credit company if there is an error. Errors can often be corrected by an eventual refund to the consumer; but if not, current U.S. law limits the consumer's liability in most cases for fraudulent transactions to $50. Substantial errors are almost always caught, and small errors, if not caught, do only minimal damage to the consumer. Financial fraud is not uncommon; but credit card companies have enormous staffs that specialize in reducing its incidence and lowering its cost; they write off the remainder as a cost of doing business.

But with i-voting, the situation is completely different. There is no way for anyone to check after the fact how anyone voted. In fact, it is important that a voter not even be able to verify that his or her own vote was recorded correctly, for that could open the door to vote coercion and vote selling, and it could also lead to a large number of almost certainly false claims that the vote reported after the fact was not what the voter thought he or she originally cast.

Without a way to check on a vote, it is difficult to detect vote fraud committed through the use of stolen authentication information or through malicious software on the voter's machine, and it is impossible to correct even if it is detected. Hence, we have no choice but to go to great lengths to prevent electronic vote fraud in the first place.

7 Internet voter education and support

No i-voting system should be fielded without a comprehensive voter education program in place to explain it to voters. At county-controlled i-voting sites there should always be someone on hand to explain to voters how they should authenticate themselves, and to offer assistance in case of any technical problems encountered during voting. It is essential that voters not be intimidated by the mechanics of i-voting, and that they have a clear mental model to use as a guide.

For home or workplace-oriented systems, there should be comprehensive documentation online, and also a "practice" site, where voters can go through the motions of i-voting, with the understanding that practice votes do not count and that they are free to experiment. Voters should be encouraged to experiment with the i-voting system, and practice the whole procedure using an alternate site before connecting to the real vote servers and casting real ballots.

Many technical problems will surely be encountered when home or workplace i-voting is first tried, and it is essential to have help resources available to guide voters through them. For example, the client software will have to work on a very wide variety of voter-owned configurations, and inevitably there will be configurations or ISPs not supported. Such situations must be handled as gracefully as possible.

The procedures for Internet voting will be unfamiliar to voters in the first few elections, and the rationale for any extra steps necessitated by security concerns will not be widely apparent, and may, in fact, be resented. Vendors of i-voting systems should also be prepared to conduct a comprehensive voter education media campaign to explain how i-voting works, and why, and that they always have the alternative of going to the polls if they encounter problems. There are many features of such a system whose purpose and functioning will not be obvious; voters will quite reasonably wonder if the system is secure. The online documentation should include answers to such potential voter questions as:

- How is it that my vote is private, when I am warned all the time that email is not very private at all?
- Why is voting more complex than buying items from an online store?
- What do I do if my computer crashes while I am voting?

Finally, vendors should be prepared with abundant technical support for voters who are having trouble during i-voting. Both telephone support and live online support are desirable, with quick enough response that voters do not abandon i-voting out of frustration.

8 General Requirements for i-voting systems:

The Internet Voting Task Force did not attempt to design a system for i-voting. Rather, we have concentrated on specifying *requirements* that such systems must meet. There are many possible implementations that will meet all of these requirements; the actual designs will reflect the influence of the Legislature, vendors, certifiers, and county procurement processes.

The following requirements recommended by this Task Force apply broadly to i-voting systems, and are not tied to any particular step in the process.

Requirement: In addition to certification with respect to traditional criteria for voting systems, any i-voting system should be certified by a Technical Review Committee composed of experts in computer- and communication security and privacy, including experts in cryptography.

The security and privacy of i-voting systems will depend critically on the entire range of computers, networking, and software used in both vote servers and vote clients, and also on careful end-to-end analysis of cryptographic authentication, and privacy protocols. The kinds of expertise sufficient to certify the traditional paper-based or mechanical voting systems are wholly inadequate for i-voting systems.

Requirement: i-voting systems should be recertified regularly.

The computing world, changes rapidly. At this stage in history the software that is commonly available for servers and clients may change considerably within any two-year election cycle, so that what was a good, efficient architecture one year may be very inefficient, or even incompatible with widely available systems, only two years later. This is particularly true of security systems and infrastructure.

Eventually cheap special-purpose voting devices may appear on the market; or perhaps more general machines with strong security architectures built-in to the operating system will become widespread (as opposed to the extremely insecure PC environments of today). Such eventualities would call for re-evaluation of i-voting systems, perhaps with the decertification of older systems and certification of new ones. (This is not as expensive as it seems; costs for all kinds of hardware will continue falling so fast that for the foreseeable future two-year-old systems will always be substantially depreciated anyway.)

Requirement: Laws against vote fraud must be reviewed in the context of i-voting.

Current laws regarding vote fraud, vote coercion, vote selling, and other election violations were drafted with paper-based balloting systems in mind. All such laws should be re-examined, and extended or

broadened where appropriate, to be sure that they prohibit interference with the privacy and security of i-voting as well.

Special consideration should be given to criminal penalties for

- Internet attacks of any kind on vote servers or other election-related computers;
- unauthorized use of encryption keys, PINs, passwords, or other authorization or authentication information, belonging either to voters or election officials, that are intended to protect the privacy and security of voting;
- deliberate interference with ballot transport on the part of Internet service providers or any other data transport companies;
- creating or knowingly distributing malicious software designed to compromise the security or privacy of i-voting;
- monitoring or interfering with the voting process, or violating ballot privacy, through systems for remote monitoring, sharing, or remote control of other computers.

Consideration should also be given to the legal recourse California may have if its i-voting processes are attacked through the Internet from foreign locations. The Federal government should consider international law or treaties to cover the case of one country's citizens interfering by Internet with the elections of another.

Requirement: The secrecy of a voter's ballot choices should be preserved, and every reasonable technical means should be used to prevent anyone from violating ballot privacy anywhere along the path from the voter to the canvass.

The natural response to this requirement is that ballots must be encrypted for transmission from the vote client to the vote servers, and we require that. But there are other potential threats to voter privacy that may occur *before* the ballot is encrypted. There are many standard commercial or freeware systems that allow one computer to monitor another, or "share" files or devices with it, or control it, through a network or through the Internet. These tools are usually quite legitimate; they are used by traveling workers who want access to the home base machine, by system administrators in the management of networks, by managers monitoring the work of employees, and often in home situations where a knowledgeable person helps maintain the computer of a less knowledgeable person, and does so remotely.

But remote monitoring or management software can also be used by a person at one computer to spy on someone who is voting at another computer, or even to control the voter's computer during voting. Voting software should therefore be designed to check for the presence of the common kinds of remote control

software, and it should then inform the voter, and not allow voting on the remotely monitored or controlled machines.

Even in cases where no known remote monitoring software is found in the configuration, i-voting software should check to see if the computer is networked at all, via LAN, or open PPP or SLIP connection, or wireless connection, or any other means; if so, it should warn the voter that it may still be possible for voting to be monitored through another computer on the network, and that the voter should not use a networked computer for voting if he or she is concerned about privacy.

In voting software configurations designed for kiosks, the configuration simply should not have any remote control or remote monitoring facilities at all.

Requirement: The ballot that is transmitted to the vote server must be an accurate copy of the voter's choices, with no reasonable possibility of undetected modification anywhere in the transmission path in any of the intervening computers and networks, including within the voter's own computer.

This requirement may sound fairly direct to people unfamiliar with computer security, but it is probably the most difficult-to-satisfy requirement in this document, and it may disqualify many otherwise attractive i-voting systems. It is vital that vendors take this requirement seriously, and that certification authorities do likewise.

Requirement: Internet voting should not continue through Election Day, i.e. there should be a time in advance of Election Day, fixed by law, when i-voting is cut off.

It is only natural that voters will wait until almost the last hour to vote by Internet. As with absentee balloting, there is an incentive for voters to wait until near the deadline so that they will have the most time to study the candidates and issues, and so they will be able to watch for as long as possible how the campaigns develop.

But with i-voting, waiting until the last minute can be risky. The first problem is that the voter's own computer might encounter hardware or software trouble. If this were to occur in the last hours of election day, such a technical problem might prevent the voter from voting because he or she will not have time either to correct it or to go to the polls.

Another concern is that a large fraction of the entire i-voting electorate can be expected to wait until the last hours to vote. Since the heaviest vote load will hit the vote servers then, it is the most likely time for overload, attack, or failure of the vote servers or of the communication links to them. Such a problem in the last hours of election day would effectively disenfranchise all procrastinating Internet voters.

For these reasons it seems wise to close i-voting a day or two before election day itself. That way, voters will be much less likely to be disenfranchised because of technical failures, either of the client machine or of the vote servers. If a voter is not out of town, he or she would be able to vote (provisionally) at the polls.

Requirement: During the i-voting window, test ballots should be regularly transmitted from all county-controlled vote clients to verify the end-to-end integrity of the entire system.

Throughout the time when i-voting is permitted, county officials should cause special test ballots to be submitted from all of the Internet vote clients under its control as part of a continuous, online logic and accuracy (L&A) test. These ballots would be indistinguishable from real ballots for all purposes except that they would not count in the final vote tally. County officials should know exactly how many test ballots are sent, and when, and from which machines, and what "votes" the test ballots contain, so that any lost ballots, extra ballots, or changed ballots can be immediately detected and appropriate action taken.

For vote clients not under county control, e.g. in homes or institutions, this procedure may not be practical. But some other L&A protocol that makes it more difficult for malicious code to interfere with voting without being detected should be employed.

9 Requirements for the Vote Server Data Center (VSDC)

The VSDC, for purposes of this document, is that part of the infrastructure that receives ballots from the Internet and secures them. It may be replicated, it may be geographically distributed, and it may or may not be at the same location as the rest of the vote-handling infrastructure. We also assume that the VSDC may be managed by a vendor or contractor to the county, rather than by county employees.

However the vote-handling infrastructure is architected, there are strong engineering requirements on the design and location of the VSDC. In the following requirements, quantitative estimates of the engineering parameters required depend strongly on the size of the county and significance of the election. The certification panel and the county procurement personnel should make sure that the actual fielded system is built to a scale appropriate to the county or counties in question.

Requirement: The VSDC must be physically secure—at least as secure against physical intrusion as the county election agency where votes are stored and tallied.

Locked doors and guards would be prudent, especially in the last days of the i-voting window.

Requirement: The VSDC must be engineered for highly reliable vote storage.

The highest priority mission of the VSDC is to store ballots (in encrypted form) and, above all, not to lose any of them. This requires that the storage system used for votes must be redundant, must be invulnerable to power failures, and perhaps make use of write-once storage, such as CD-R.

Requirement: The VSDC must be architected for high availability.

"High availability" means that the VSDC must be up and available for voting for all but a negligible fraction of the time during the window in which i-voting is permitted. It should be engineered with redundant servers, redundant communication, and with smooth failover procedures so that if one resource goes down, the others remaining can automatically take up its slack with no loss of votes and minimal disruption.

Figure 1 shows, for example, redundant vote servers, each with redundant disks, and redundant connections to the Internet through multiple ISP's. Redundant resources should be architected for smooth failover. The VSDC will also need a battery-powered UPS (uninterruptable power supply) and a backup power generator to guard against power failures.

Requirement: The VSDC must have sufficiently high-bandwidth connections to the Internet.

It will need enough communication capacity to handle the maximum rate of votes that might reasonably be expected in the last hours that i-voting is permitted, and do so even if some of the connections to the Internet are down or are under denial-of-service attack.

Requirement: The VSDC servers must have sufficient computational performance to provide responses back to voters in a few seconds.

Fast response indicating that their ballot has been received is important for voter satisfaction and confidence, and it must be achieved even if some of the vote servers are down.

Requirement: The VSDC should have a connection to the county premises if it is not located there.

The connection does not need to be as secure, high-performance, or highly-available as the other parts of the VSDC.

Requirement: The VSDC must be equipped with systems and procedures to withstand most attacks on its servers, including denial-of-service attacks.

This requirement is generally met partly with some kind of "firewall", a system of special computers that filter traffic, and partly through vigilance on the part of operators, who should be wary of attacks and prepared to take fast action.

The firewall should block all incoming packets on all ports except those involved in voting, and should be configured to filter malformed packets and any other suspicious traffic.

A denial-of-service attack on a server is an attack designed either to clog the communications channels leading to the server so that requests to it and responses from it cannot get through, or to crash the server repeatedly so it gets no work done, or to overload the server with fraudulent requests that force it to take all of its time checking and rejecting them instead of dealing with legitimate requests. Such an attack does not aim to take control of the server or get it to do any specific thing; it just aims to keep the server from getting its work done, thereby "denying service" to all users as if there were a massive system failure. In the case of the vote servers of the VSDC, a successful attack would effectively prevent it from accepting votes.

There are numerous well-known denial-of-service attacks. Many can be ameliorated by careful firewall configuration. Others can be defended with the help of excess resources on the server, and redundant servers with smooth failover techniques. But the most comprehensive approach is to vigilantly monitor the server(s) and networks for such an attack and to be prepared quickly to cut communications with the network(s) from which the attack originates (although that would also cut off voters originating from that

network). This requires skilled systems personnel. Any vendor or contractor who bids on a contract for i-voting in a California county should demonstrate that they have the resources and skills needed to defend against such attacks.

10 Requirements for the Internet Voting Process

The following sections list detailed requirements for each step of the i-voting process, more or less in the order they occur from the perspective of a single voter.

i) Request for Internet balloting

Requirement: Voters must request i-voting in writing with an original signature; they must re-request for each new election, and must not request both an absentee ballot and i-voting in any one election.

Voters who wish to vote via the Internet must request it in writing, with an original hand-written signature, in a manner and under rules essentially the same as for requesting an absentee ballot in California. The two requests could be on the same form, with a check box indicating which the voter wants. A signed, written request for i-voting is essential, because comparison with the signature on file with the county registrar of voters is the only test there is in the current system that the requestor is eligible to vote. If other forms of voter authentication, such as thumb print, driver's license number, or digital signature are ever added to the requirements for voter registration, then this requirement for hand signature on the request for i-voting, or even the requirement for the request itself, can be changed accordingly.

It is absolutely essential that all signatures on requests for i-voting be checked against the signature in the registration file before issuing authorization for i-voting. Unlike absentee ballots, which will be accompanied by another original hand signature that can be checked before counting, Internet votes will have no hand signature; hence checking the signature on the request for i-voting is mandatory.

In accordance with California absentee balloting procedures, voters should not be permitted to request i-voting permanently (with the exception of voters with medical need, or voters living in rural precincts where there are no polling places), for the same reason that they cannot normally request to vote by absentee ballot permanently—it is too easy for Internet ballot authorization to be issue automatically over and over, long after the voter has moved away or died. Furthermore, the procedures for requesting

absentee ballots, or the county's response, may change in the first few elections in which i-voting is tried, so widespread permanent i-voting authorization may become a burden to administer.

Voters should not be issued both authorization for i-voting and an absentee ballot, even if they intend to use only one or the other. The verification that they have not double or triple voted (by also showing up at the polls) is too much of a clerical burden on election staffs.

ii) Authorization for Internet ballot

Requirement: The authorization for Internet balloting can be in various forms depending on the design of the i-voting system as a whole. But any authorization must provide a way of linking the eventual vote cast using that registration to the registration record for that voter, so that it can be determined beyond a reasonable doubt that each Internet vote is associated with a registered voter in the proper district, and that at most one vote is counted for any voter, whether at the polls, or by absentee ballot, or by Internet voting.

A county's response to the request for an Internet ballot will normally be to issue an *authorization* for Internet balloting to the voter who requested it. The authorization will be some combination of cryptographic keys, or PINs, or both, possibly accompanied by voting software. The authorization may be handed to or mailed to the voter on computer readable media, or it may be emailed to the voter, or it may be made available password-protected by a randomly-generated password over the Web; different i-voting systems may differ on this point.

The fact that a voter has been authorized for i-voting, and any security information associated with it, must be stored by the county for use in authenticating the ballot and preventing double voting later. It must be possible to cancel a voter's authorization in case of it is lost or compromised in some way.

iii) Loss of Internet ballot authorization

Requirement: Any system must be able to handle the voter's loss of, or failure to use, authorization for Internet balloting.

If a voter loses, or claims to lose, his/her Internet ballot authorization, or if that authorization for some reason fails to work to allow voting, then the voter can request a new Internet authorization, or an absentee ballot. Before either such request is granted, the old authorization must be canceled. The voter may

instead just go to the polling place on election day and vote with a provisional ballot even if his authorization for i-voting has not yet been canceled by the county.

iv) Voter authenticates himself/herself

Requirement: Voters should be provided with an authentication code from the county that is combined with a personal identification number (P.I.N.) that will allow the voter to authenticate him/herself for the I-voting system.

No single interception of an "out-of-band" transmission should allow an individual to cast a fraudulent ballot. Voter authentication codes provided by the counties can be combined with a number or password requested by the voter to ensure that at least the same level of security that is achieved in the absentee ballot process is available for Internet ballot. In paper absentee ballots, the theft or interception of a blank ballot would not necessarily result in the successful voting of an illegal ballot because the voter is required to affix is or her signature to the exterior ballot envelope. That same level of security should be mirrored in Internet voting.

v) Voter brings Internet ballot to screen

Requirement: The screen on which the user views the ballot must be capable of rendering an image of the ballot in any of the languages and orthographies required by law for paper ballots.

Today, federal law requires some California counties to print ballots in English, Spanish, Tagalog, Vietnamese, Japanese, and Chinese. Counties can add to this list; Los Angeles County, for example, includes Korean.

Requirement: No contest, either for an office or a proposition, should be split across two screen pages.

If there are six candidates for an office, then all six should be visible on a single screen page in order not to disadvantage candidates at the bottom of the list. For systems employing voting devices having displays other that those used for PCs, this puts a constraint on how small the screen should be.

Requirements: The application used for voting should not display or play any advertising or commercial or logos of any kind, whether public service, commercial, or political.

Web browsers and similar programs are capable of displaying text, graphic, audio, animation, and video advertising. Many times the ads are inserted by the providers of a Web site; sometimes they are added by another "framing" site; still other times they are inserted by the Internet service provider. To be consistent with the principle behind the law that there should be no advertising or campaigning within a certain radius of the polling place, we recommend that there should be no advertising in the "window" that contains the voter's ballot, or popped up as a result of retrieving the ballot. The ballot must not have the appearance of being "sponsored by" any person or organization. This requirement may have no simple technical solution, and may thus have to be backed up by law.

However, this does not mean that voters cannot have political information and advertising in *other* independent windows at the same time they are viewing the ballot. Just as people are permitted to take any material they wish into the voting booth, there is no reason why they should not be able to visit other web sites, including political sites, while voting (as long as other security requirements are met, e.g. no ActiveX controls, JavaScript scripts, Java Applets, etc.).

Requirement: Multi-page ballots should be easily navigable by voters, with no way to get lost or leave the balloting process except deliberately.

If the ballot is in the form of a Web page it should contain no hyperlinks to other sites, which would be distracting, and might cause voters to get lost while voting.

vi) Voter makes choices

Requirement: Over-voting (voting for more candidates than permitted for a single office) must be prevented.

The voter should be notified, as soon as the he or she attempts to vote for too many candidates, and no ballot with over-voting should be transmitted to the server. This service to voters is similar to that provided in some other voting systems, e.g. mechanical voting machines and some mark-sense balloting systems.

Requirement: Voters should be able to point and click to make their voting selections, or type a write-in name. They should be able to navigate back and forth within the ballot to change selections freely until the moment when they click the final button that irrevocably transmits their ballot.

A smooth, easily understandable, navigable, and fairly platform-independent human interface is vital to voter acceptance.

Requirement: Needs of voters with disabilities or impairments should be accommodated.

It should be possible for an audio version of the ballot to be read by the computer to the sight-impaired, and the position of the screen and keyboard/mouse (or other input device), should accommodate wheelchair-bound voters.

Requirement: Voters should be able to type write-in candidates' names in any language or orthography required by law for paper ballots.

Internet voting should be as accessible to non-English speakers as it is to English speakers, just as is true for paper ballots.

Requirement: The actual contents of the voter's votes on the client computer should be kept only in volatile memory, if possible, so that it will be automatically erased in the event of a power failure or rebooting. Votes should not be written to long-term storage on the client machine or for any reason, even in encrypted form.

A voter's vote should not be stored in a file on the client machine, even a temporary file, and it should not be paged out to secondary storage as a result of virtual memory. It also must not find its way into any log, cache, index, cookie, or any other long-term record. And since the encryption key(s) used in encrypting the vote may be stored in or near the voter's computer, this extends even to encrypted votes.

vii) Voter casts ballot

Requirement: No vote must be transmitted before the voter clicks on a next-to-final button labeled, for example, "Send Ballot". After clicking, the voter must be told that sending the ballot is

irrevocable and must be asked to confirm his or her intention to send the ballot by clicking a "Confirm" button. If the voter does not then click the "Confirm" button, he or she should be able to return to the ballot to continue voting; but if he or she does, then voting is complete.

It is important that the voter not accidentally send the ballot prematurely, because there can be no way to retrieve it, complete it, or vote again, and the voter would then be at least partially disenfranchised.

Requirement: Immediately after the ballot is sent to the vote server, and without waiting for feedback from the server, or immediately after the voter clicks on the "cancel" button, all record of the vote must be deliberately erased from the voter's computer.

Any choices the voter made should first be erased from the screen. Also, the voter's choices are presumably held unencrypted in the computer's RAM, and would remain so indefinitely unless the voting application deliberately zero's them. (Memory deallocation is not sufficient.) If the voter walks away from the computer after voting, it must be infeasible for someone else to walk up to it and apply any software tool to recover the votes. If feedback from the vote server indicates that the vote was not accepted, and the voter wants to try again to vote by Internet, he or she must start over.

viii) Ballot transmitted to vote server

Requirement: The ballot, along with a timestamp, voter's identification, precinct, and any other appropriate information, must be transmitted to the vote server in encrypted form to protect the privacy and integrity of the information.

It must be infeasible for anyone who taps the communication links between the voter's computer and the vote server to read the ballot, or any of the associated information, or to tamper with any of it in a way that might go undetected. It must also be infeasible to inject a duplicate of the encrypted ballot and have that counted as an additional vote.

ix) Vote server receives ballot

Requirement: The ballot transaction is atomic. A ballot must be either wholly accepted, or wholly not accepted, by the vote server. There must be no middle ground.

If it is accepted, the voter should not be able to vote again; if it is not accepted (including the case of not being received), the voter is permitted to vote again, either by Internet or at the polls by provisional ballot.

Requirement: The vote server that receives a ballot should immediately check it to ensure that it is formatted correctly. If it is, the vote server should immediately store the ballot, still encrypted, on a permanent medium (e.g. a CD-R disk) so that any subsequent power or equipment failure will not lose the ballot.

If the check of the ballot fails, the voter should be notified and given advice about what to do, i.e. try again, or give up and vote at the polls. In either case, valid or not, the vote server should store the vote permanently and redundantly for later decryption and canvass. The encrypted ballot, valid or not, may be considered part of the audit trail in case a recount is called for, or the election is challenged in court.

Requirement: If the vote servers are managed by contractors, rather than by election officials, then no keys or other tools for decrypting ballots should reside on the vote servers or be available to the contractors.

All such keys must remain strictly in the hands of election officials.

x) Vote server sends feedback to voter's screen

Requirement: Within a few seconds of receiving the ballot, the vote server should attempt to notify the voter of whether or not the vote was successfully accepted.

When the voter is finished, i.e. any time after hitting the "confirm" or "cancel" button (even if feedback from the server has not arrived) then the voter and should be able to just walk away without "closing" or "shutting down" anything, and still be guaranteed the privacy of the vote. If the vote was not accepted, then the voter may start over, or may vote by provisional ballot at the polls.

Requirement: If no feedback comes back to the voter's computer within a reasonable time, for any reason, then the voter is entitled to assume that the vote was not accepted, and may try again to vote by Internet, or may vote by provisional ballot at the polls.

There are many reasons why the feedback might not arrive at the voter's computer. Computer failures, software crashes, or communication failures, either at the vote server, or at the client, or in the Internet infrastructure in between, are all capable of preventing the ballot from being delivered to the vote server, or preventing the feedback from being delivered back to the voter. Most of these cases are completely out of control of the voter, and are all indistinguishable from his point of view. In particular, the voter cannot tell, in the absence of feedback, whether the vote was rejected for some reason, or was accepted but the feedback was lost. So the voter should be entitled to vote again.

If the vote in fact did arrive and was accepted, but the feedback was lost, then the fact that the voter votes a second time, either by Internet or by provisional ballot, must be detected, and the second (and subsequent) ballots excluded from the canvass. Double voting, in this case, should not be held against the voter. Since the two ballots need not agree in all contests, there needs to be a strict rule about which one takes precedence, and the choosing the first one is the most reasonable; choosing the second one would be tantamount to allowing the voter to change his or her vote.

xi) Voter can ask for confirmation that he/she voted

Requirement: There must be a mechanism that voters can use to determine the status of their vote, i.e. whether or not it has been accepted and authenticated.

Voters should also be able to authenticate themselves online and then query whether or not their vote has been accepted and authenticated. The original feedback a voter receives only indicates, if positive, that their vote was *accepted*, i.e. stored securely. But, depending on the voting protocols, it may be that the vote is *authenticated* only later.

In order for voters to be confident that their Internet vote will be counted in the election, and that they do not have to vote again, there must be a mechanism for voters to query whether their ballot was accepted and authenticated. They may want to check that it was accepted in case the acceptance feedback did not get to them for some reason when they tried to vote. And they may want to know that it was later authenticated so that they need not go to the polls to cast a provisional ballot.

Note that this requirement goes slightly beyond what is possible for current absentee ballots.

Requirement: After the voter has sent the ballot to the vote server, there must be no way for anyone, even the voter, to determine *how* he or she voted in any contest. In particular, there must be no way that a voter can prove to a third party how he or she voted.

Because of the danger that voters might be coerced or paid to vote a certain way, it is important that voters have no way of proving after the fact how they voted, even voluntarily.

Of course, it is possible that someone might be watching over the shoulder of a voter while he or she is filling out an Internet ballot, and no technical requirement can prevent that. But such a possibility applies also to someone filling out a paper absentee ballot as well, so i-voting is no less private.

xii) Votes transmitted from vote server to canvassing machines

Requirement: Internet Voting systems must be capable of accurately tabulating the results and integrating the results with the county's primary voting system.

xiii) Authentication of votes and separation from voter identification

Requirement: The county election system must be able to verify the authenticity of a ballot before the votes on the ballot are viewed or counted.

Similar to a paper absentee ballot, Internet ballots should be verified for authenticity before the authenticating information is stripped from the ballot. The verification of the authenticity of the ballot should ensure the true source of the message. This must ensure that an electronic ballot really is from the person it claims to come from, and not just from someone trying to electronically impersonate that person.

As in the paper absentee ballot process, once the ballot is separated from the authenticating information on the envelope, the ballot must be incapable of being traced to the voter who cast it.

The voted ballots are decrypted and counted after the authenticating information is reviewed and removed from the ballot.

xiv) Canvassing of votes

Requirement: The Internet voting system must be capable of accurately tabulating the results of all ballots cast. The canvass should only be conducted after the close of polls on election day.

xv) Maintenance of auditing information

Requirement: Decrypted ballots must retained in a secure format to allow for subsequent auditing and recount procedures.

xvi) *Human security*

Requirement: In accord with the rules for handling absentee ballots, no single election official should be able to delete, change, forge, or violate the privacy of Internet ballots.

Election officials are bound by rules and procedures governing the handling of ballots that are designed to ensure that the privacy of votes is respected, that no ballot is lost or unaccounted for, and that no single employee can change, forge, or destroy a ballot. Absentee ballots, for example, are always handled in the presence of at least two employees. Ballot envelopes are face down so that the signature on the ballot envelope is not visible when the ballot is separated from the envelope. And all absentee ballots mailed out are coded and accounted for, even if they are not returned by the voter.

Analogous procedures are also necessary for "handling" Internet ballots. Internet ballots will be held in files and operated upon by software tools for validation, for separating voter identification from votes, and for canvassing. Any i-voting system must have security mechanisms in place that guarantee at that at least 2 employees should concur whenever any critical operation regarding the processing of Internet ballots takes place, i.e. the passwords or cryptographic keys of at least 2 employees are required to operate on votes.

11 Glossary

ActiveX control: A program packaged in a format designed by Microsoft that is downloaded from a web server to a client browser and run within the browser, all as a mere side effect of visiting a web page.

Applet: A program in Sun Microsystems' Java programming language that is downloaded from a web server to a browser and run in the browser as a side effect of visiting a web page.

Atomic: A multi-step operation is atomic if, whenever it is attempted, it either fails completely, accomplishing nothing at all, or succeeds completely, accomplishing all of the steps, but never stops in an intermediate, partially-completed state.

Authentication: Verification of the true source of a message. In the case of i-voting, this refers to verification that an electronic ballot really is from the person it claims to come from, and not just from someone trying to electronically impersonate that person.

Biometric: A digitizable characteristic of a person's physiology or behavior that uniquely identifies him or her. Examples include thumb print, DNA sample, voice print, hand-writing analysis, etc.

Browser: An application program such as Microsoft Internet Explorer or Netscape Navigator that allows the user to navigate the World Wide Web, and interact with pages from it.

Certification: The process the state uses to determine that a voting system meets the requirements of the California Election Code and can be used by any county that decides to select it.

Client: In a common two-computer interaction pattern, one of them, the *client*, initiates a request, and the other, the server, acts on that request and replies back to the *client*. In the case of i-voting, "client" refers to the voter's computer that initiates the process of voting, and the server is the computer that accepts the ballot and replies to the client that it accepted it.

Cryptography: The mathematical theory of secret codes and related security issues.

Decryption: Decoding an encrypted message (usually using a secret key).

Digital signature: Cryptographically-generated data block appended to a document to prove the document was processed by the person whose secret key was used to generate the data block.

Encryption: Encoding (i.e. scrambling) a message using a secret key so that anyone intercepting the message but not in possession of the key cannot understand it..

Failure tolerance: The ability of a system to continue to function in spite of the failure of some of its parts.

eCommerce: Electronic commerce, i.e. financial transactions conducted over a computer network or the Internet.

Email: Electronic mail, i.e. messages and documents sent from one party to other specific, named parties.

Firewall: One or more computers standing between a network ("inside") and the rest of the Internet (outside). It intercepts all traffic in both directions, forwarding only the benign part (where "benignness" may be defined by a complex policy), thereby protecting the inside from attacks from the outside.

HTML: Hypertext Markup Language, the notation used for formatting text and multimedia content on web pages.

HTTP: Hypertext Transfer Protocol, the communication protocol used between web browsers and web servers for transporting web pages through the Internet.

i-voting: Internet voting

Integrity: Protecting data from undetected modification by unauthorized persons, usually through use of a cryptographic hash or digital signature.

Internet: The worldwide system of separately-owned and administered networks that cooperate to allow digital communication among the world's computers.

IP: Internet Protocol, the basic packet-exchange protocol of the Internet. All other Internet protocols, including HTTP (the Web) and SMTP (email) use it.

IP Address: A unique number (address) assigned to every computer on the Internet, including home computers temporarily connected to the Internet.

ISP: Internet Service Provider; a company whose business is to sell access to the Internet, usually through phone lines or CATV cable, to homes, businesses, and institutions.

Key: A typically (but not always) secret number that is long enough and random-looking enough to be unguessable; used for encrypting or decrypting messages.

Key pair: A pair of keys, one used for encrypting messages and the other for decrypting them. Used in public key cryptographic protocols for authentication, digital signatures, and other security purposes.

Kiosk: A booth- or lectern-like system with a screen, keyboard, and mouse mounted so they are available to users, but with a tamper-proof computer inside and a secure Internet connection to the server.

Mirroring: Keeping two or more memory systems or computers identical at all times, so that if one fails the other can continue without any disruption of service.

LAN: Local Area Network; a short-range (building-size) network with a common administration and with a only small number of hosts (computers) attached. The hosts are considered to be sufficiently cooperative that only light security precautions are required.

Malicious code: A program with undesirable behavior that operates secretly or invisibly, or is disguised as part of a larger useful program; in this document, the same as "Trojan horse".

NC: network computer; a widely-discussed hypothetical product that does not store software or files locally, but works only through a network.

Online: Generally, a synonym for "on the Internet", or sometimes, more specifically, "on the web".

Out-of-band communication: Communication through some means other than the primary channel under discussion. If the primary communication channel is the Internet, then out-of-band channel might be via U.S. mail, or a voice telephone connection, or any other channel that does not involve the Internet.

Packet: The smallest unit of data (along with overhead bytes) transmitted over the Internet in the IP protocol.

PC: Personal computer; any commercial computers marketed to consumers for home or business use by one person at a time. In 1999, this includes Intel-based computers (and clones) running a Microsoft operating system or a competitor (e.g. Linux, BeOS, etc.), and it also includes Macintoshes.

Plug-in: A software module that permanently extends the capability of a web browser.

Privacy: Protecting data from being read by unauthorized persons, generally by encrypting it using a secret key.

Private key: A key, or one member of a key pair, that must be kept secret by one or all members of a group of communicating parties.

Protocol: An algorithm or program involving two or more communicating computers.

Public key: One member of a key pair that is made public.

Public key cryptosystem: A cryptographic protocol involving a pair of keys, one of which is made public and the other held secret.

Redundancy: Excess storage, communication capacity, computational capacity, or data, that allows a task to be accomplished even in the event of some failures or data loss.

Replication: A simple form of redundancy; duplication, triplication, etc. of resources or data to permit detection of failures or to allow successful completion of a task in spite of failures.

Script: In the context of this document this term refers to a program written in the JavaScript language, embedded in a web page, and executed in browser of the web client machine when it visits the web page.

Security: General term covering issues such as privacy, integrity, authentication, etc.

Server: In a two-computer interaction pattern, one of them, called the client, initiates a request, and the other, the *server*, acts on that request and replies to the client. In the case of i-voting the computer that receives and stored the ballots from voters is the server.

Spoof: To pretend, usually through a network, to be someone or somewhere other than who or where you really are

Trojan horse: A program with undesirable behavior that operates secretly or invisibly, or is disguised as part of a larger useful program; in this document, the same as “malicious code”.

Tunnel: A cryptographic technique in which a computer is in effect attached to a remote LAN via the Internet, even if there is an intervening firewall.

URL: Uniform Resource Locator, i.e. a name for a web page, such as http://www.vote2000.ss.ca.gov .

USB port: Universal Serial Bus port; a port (connector) on newer computers used for high speed serial communication with attached devices.

Virus: A Trojan Horse program that actively makes, and covertly distributes, copies of itself.

Vote client: The computer that voters use to cast their ballots, which are then sent to the vote server.

Vote server: The computer(s) under control of the county that receives and stores votes transmitted by Internet from vote clients.

Web: The world-wide web, or WWW; the worldwide multimedia and hypertext system that, along with email, is the most familiar service on the Internet.

Web site: A collection of related web pages, generally all located on the same computer and reachable from a single top-level "home page".

Web page: A single "page" of material from a web site.