Testimony for the U.S. Election Assistance Commission
Neil McClure, Vice President
Hart InterCivic, Inc.
May 5, 2004

Thank you very much for the opportunity to present testimony before the Commission on behalf of Hart InterCivic.

Hart InterCivic traces its roots to 1912 as a provider of products and services for state and local government. For over 90 years, Hart has been supplying election products and services, including voting supplies and paper balloting products, to election customers across the nation. Throughout our company's history, Hart InterCivic has supported all types of elections from paper, to lever machines, punch cards, optical scan and now Direct Record Electronic (DRE) elections.

In the 1980s and 1990s Hart was a reseller of optical scan voting systems. Recognizing that elections, like many other essential government functions, were evolving toward electronic solutions in place of paper, we decided to expand our company's product line to include electronic voting systems. Our initial goal was to address two niche markets: (1) systems to support voters with disabilities and (2) systems to help local election officials manage early voting in person. In the latter case, electronic systems are ideal to eliminate the logistical complexity of paper ballots in multi-precinct, possibly multilingual polling places.

However, the architecture of our system was designed to accommodate *all* user requirements, with a major focus on security. With our experience in the elections industry, we understood the importance of the user requirements, product quality, product reliability and above all else, ensuring the accuracy and integrity of election data.

In the summer of 2000, following a three-year development effort, our DRE system, the eSlate™ Electronic Voting System, was federally certified by the National Association of State Election Directors (NASED) and introduced to the marketplace. The system was first used in live elections during the 2000 General Election.

Since the November 2000 elections, the eSlate System is now installed in 28 jurisdictions in 9 states including Texas, California, Colorado and Virginia. The eSlate System has been implemented in Harris County, Texas (Houston) and Orange County, California – respectively the 3rd and 5th largest counties in the U.S., and the two largest counties to have purchased and successfully implemented electronic voting systems.

Since the introduction of the eSlate System, we have released new eSlate System applications to support storage and warehouse management of the equipment, distributed collection of cast vote records, candidate rotation, and multiple language support. These features represent the focus of our development resources over the past three years and the priorities of this development have been determined by market and customer requirements. The market has determined that the next feature required of all electronic voting systems is increased security. Therefore, we are focusing on accelerated implementation of additional, enhanced security features, many which were already a part of our original architecture and product roadmap.

Our recognition that security is an important feature of an electronic voting device is well documented in the disclosure of U.S. Patent Number 6,250,548, filed in 1997, for which I was the principal author (the McClure ‘548 patent). For example, column 16, line 58 of this patent references “executing public encryption algorithms” to protect portable memory devices storing election data. While the McClure ‘548 patent does not provide a complete treatment of voting device security, it does point out a fundamental awareness of security requirements in the system design. When we introduced the eSlate System, security requirements primarily addressed voter fraud and the protection of voter privacy and anonymity.

The electronic voting device market now requires higher levels of security, but this has been mistakenly identified by some as a need for paper ballot receipts. The introduction of the paper ballot receipt concept has been proposed to mitigate only one risk to DREs -- whether the DRE actually records the vote as reported to the voter on the voting device’s display screen or whether it is changed or corrupted as a result of a malicious attack. There are other security risks to DREs, but this is one that critics of electronic voting systems identify as being mitigated by the adoption of a paper ballot receipt.

Opponents of electronic voting further propose that any additional risks posed by DREs, or risks of fraudulent totals following the voting process, would also be mitigated because of the existence of the paper ballot receipt. Paper ballot receipts are subject to tampering and while there is a great deal of history to support the susceptibility of paper ballots, I will not address this issue here today.

The paper ballot receipt is a solution to a problem, and standards should not prescribe solutions but rather identify requirements to prevent problems. So what is the problem? As I mentioned previously, the paper ballot receipt is being proposed to mitigate the risk that the voting device may not record the vote as cast, and that voters and election officials would have no way of knowing if this occurred. So the problem is a matter of trust, and today there is the perception that DREs cannot be completely trusted. Therefore, to address this problem, the requirement is for the DRE to be "trusted," not add a paper ballot receipt. Once a reasonable level of trust is determined, and if this trust level cannot be achieved by a particular DRE, then the DRE must be treated as an un-trusted element and additional mitigation requirements would apply, such as the addition of a paper ballot receipt.

Trust must be established in relationship to the threats to a system. Any real security analysis begins with identifying the threats to a system or device and then performing a risk assessment where a threat is attributed a probability and/or likelihood that the threat would materialize into a successful attack. Many publications have identified possible threats to a DRE, but calling for a requirement of a paper ballot receipt represents a significant leap and I have not seen any published information related to the probability or likelihood of a successful attack.

To outline the complexity involved, I would like to describe the threat as I understand it and then outline the complexities to overcome in order for the threat to become a successful attack on a DRE. I will start with some reasonable premises:

- The DRE manufacturer is not engaging in a massive conspiracy to commit election fraud,

- All testing of the DRE hardware/software shows accurate recording and reporting of results,
- The attacker, or attackers, desire to remain anonymous, and
- The attack must involve more than one DRE from the same manufacturer.

In order for the threat to materialize into a successful attack, altering cast vote records as they are recorded for some desired outcome, malicious software code, specific to a particular DRE, must be written that shows the voter one outcome while recording another. From a software development standpoint, this may not be a difficult task but it does require a certain level of software engineering skill and knowledge of the electronics employed by the DRE. This provides some limitation on the number of people capable of performing such an attack.

However, in order for the attacker to remain anonymous and/or perpetrate the attack, the malicious code must be added to the existing DRE code and remain undetected or hidden. This increases the required software engineering skill level and hardware knowledge required, further reducing the field of possible attackers.

The attacker must also have access to a DRE for an extended period of time and the financial means to support development of the malicious code. For an attacker to gain access to a DRE, the attacker must either steal a device, work for a jurisdiction where the DRE is used or be an employee of the DRE manufacturer. The latter is typically referred to as a “rogue programmer.” Furthermore, the attacker must be in a position to load the malicious code onto more that one DRE sometime prior to its use in an election. This would likely require more than one person.

However, the loss of anonymity increases as the number of people involved increases. Finally, the attacker must be motivated to engage in performing the attack. These conditions limit the population of potential attackers.

As a point of contrast, the available population for a DRE attack should be compared to the population available to attack a paper-based system.

Assuming that an attacker exists, I'd like to provide an overview of the complexity involved that the malicious code, or more popularly known as a "Trojan Horse," would need to overcome. As an example, the November 2004 General Election in Orange County, California will be conducted on the eSlate Electronic Voting System and the election definition includes the following parameters:

- 2,200 precincts
- 5 languages
- 1,723 polling locations
- 91 cities and special districts

This will result in thousands of different ballot styles in the electronic ballot definition, not including the differences in presentation created by languages or ballot rotation. While the presidential race appears as the first contest, the order of the candidates and the language used can change for different presentations of the ballot, requiring the malicious code to possess control over these functions. Malicious code will need to take this enormous complexity into account. Real election are much more complex than the samples created by academics to "demonstrate" system vulnerabilities.

A common claim is that the malicious code need only search for party name and make decisions on vote alteration based on this criteria. In the case of the eSlate, the system uses a highly referenced data structure such that the words “Republican” or “Democrat” only appear once in the ballot definition file. When instructed by ballot style to display the party name, the party name is converted to a graphic image and positioned on the voting device’s screen for display. All information on the screen is in a graphical format known only by the system as a bitmap image. The point here is that the selection displayed to the voter has no reference to the party affiliation in the software code used to display the image to the voter. Again, while it would be possible to decode this process, the malicious code becomes much more complex, thereby increasing its size. As the size of the code increases, the ability to hide it decreases.

This is not meant to be an exhaustive treatment of this particular threat or risk assessment, but merely an illustration of the information that is lacking in the paper ballot receipt debate: probability and likelihood of a successful attack. Yes, it is possible, but what is the probability or likelihood of success? Probability and likelihood of success have yet to be seriously considered in this DRE security debate.

I believe that it is possible to define an evolving level of trust for DREs that is acceptable in the face of the threats that exist today and into the future. As previously described, an attacker would have to be highly skilled, well organized and a close insider to the election process. DREs are relatively new and it will take time for the hacker community to get organized enough to be in a position to mount a successful attack. We should use that time to raise the level of security required by DREs. Remember, it has taken over a decade for viruses to attack the Internet. While as a society we have accelerated the attention of the opportunity for hackers through the press and such academic programs as Rice University's "Hack-A-Vote Voting (in) Security Project." However, there remains time for a reasonable, measured response from the election community.

As an example of a first step, it would be reasonable to establish a requirement that all DREs become compliant with Level 2 of the FIPS 140-2 standard (Federal Information Processing Standard 140-2 - Security Requirements for Cryptographic Modules) by January 1, 2006. This is a federal standard that specifies the level of security for cryptographic modules and is an achievable requirement for DREs. If a DRE is not able to be compliant by this date, then it must be treated as an un-trusted device and provide alternative means to mitigate risk. During this time, the election and technical communities can be working on the next level of security requirements for DREs. Threats continually mature and become more sophisticated, which is why security is never complete. This suggested approach puts the election industry on the continuous improvement security path and maintains the benefits of DREs for election processes and the voting population.

There are several additional topics that should be addressed as part of a security standard to improve the state of DREs, but first I'd like to comment on the proposed process being recommended for the paper ballot receipt, should it be implemented. The recommendation for the paper ballot receipt proposes to provide the voter with the ability to verify that the information printed on the paper accurately reflects the selections made on the DRE. It has been further recommended that if the paper verification does not match, the voter must be given the opportunity to reject the paper verification and be allowed to vote again, up to as many as three times. This process amounts to giving the voter three opportunities to change his/her mind. The reality is, if the paper ballot does not match what the voter entered on the DRE, the system is not functioning properly and voting should immediately cease. In fact, all equipment becomes suspect at that moment and the entire election should be stopped and the appropriate legal authorities notified. Then a decision should be made whether to shut down all equipment nationally that is the same make and model or that is running on the same version of software. This is why if any form of the paper ballot receipt is implemented, a law providing for severe criminal penalties should go into effect simultaneously if a false claim is made concerning the accuracy of the system. Those of us who have experience with voters know this will occur.

This discussion leads to another area that should be addressed regarding security. There have been many reports in the media concerning irregular behavior of DREs. The DRE critics and proponents of paper ballots have pointed to these instances as further proof of their case for a paper ballot receipt. But without exception, these instances had nothing to do with security. The irregular behavior can be solely attributed to poor product quality.

This points to the need to raise the quality standards within the election industry. This can be accomplished primarily through more aggressive physical testing and a great use of software stress testing through volume simulations. The EAC should consider making national and international methods and standards for quality management and system testing a requirement for all companies developing voting system software and hardware to insure that quality processes are incorporated and audited as part of software development and voting hardware manufacturing.

In fact, simulation should become an ingrained component of system functionality and verified during ITA certification to pave the way for its use as a pre-election test, replacing the paper concept of a Logic And Accuracy Test (LAT). The practice of a pre- and/or post-election test remains valid but by applying the processes developed for paper systems to DREs a complicated, cumbersome procedure that is highly susceptible to human error results. New election test practices need to be defined that are appropriate for DREs, involving verification and validation of data and simulations that test the data path of each component of the system, from ballot definition to election reporting.

Another inappropriate paper practice that has been applied to DREs is the concept of a recount. One of the claims of the paper ballot receipt proponents is that there is nothing to recount when using DREs. However, they are applying a paper concept to an electronic device. If we look at the purpose of a recount, it is to validate and verify the outcome of an election. Marks or hole punches on paper or cards may have been misread in the initial count, and so the recount must examine the face of the ballot and ensure the voter's intent is correctly interpreted. Paper ballot counting machines require calibration and other mechanical challenges, so verifying that the tabulators are within tolerance and validating the initial count is an appropriate process, i.e. system verification. However, a recount is much more than running the paper back through the counting machines; it involves an audit of the entire election process. Historically, when an outcome of a contest has been changed due to a recount, it is not because the count from a tabulator is different, but because the results from the audit of the processes surrounding the tabulator uncovered an irregularity.

The complaint about DREs is that if a recount is done, the exact same result is reported. But that’s the point! If the same result is not reported, then there is a problem. For example, if the much discussed Trojan Horse has worked its malicious magic during a specific time window on Election Day, then an electronic recount in which data is freshly read and tabulated WILL reveal a problem. There are other steps that can and should be taken for system verification, for example, review of the audit logs along with the traditional audit of the entire election process. This reasoning is only valid if the DRE is a trusted device.

Treatment of election records also needs to be updated for DREs as well. Requiring modern electronic data storage and retrieval practices of DREs will simplify the process for election practitioners and ensure a reliable and secure retention process for the safekeeping of election evidence, should it be required for an investigation.

Another topic for consideration is the adoption of a standard for a publication of election data. Such a standard will bring new efficiencies to the election industry and provide greater visibility for public monitoring and reporting. The IEEE P1622 committee is working on this effort and your support and involvement will be key to the committee’s success.

To conclude my comments, adopting an evolutionary approach to security and addressing the necessary processes supporting DRE elections will allow secure, reliable and trustworthy elections to be conducted using electronic systems.   We should move forward with electronic voting in a deliberate and reasonable manner, celebrate the efficiencies and enfranchisement of all voters and appropriately manage the risk.

Thank you for your time.

eSlate is a trademark of Hart InterCivic, Inc.