



Premier Technical Response to the Ohio EVEREST Report

March 2008

Introduction

In late 2007, Ohio Secretary of State Jennifer Brunner commissioned a study of Ohio's voting systems called EVEREST, Evaluation and Validation of Election-Related Equipment, Standards and Testing. As stated in the 86-page executive summary, Secretary Brunner sought to provide an assessment of Ohio's voting systems that was comprehensive, independent, balanced, and objective.

Premier notes the following observations:

- The study does not, in fact, constitute a complete and comprehensive review. Ohio's voting system, like that of any state or jurisdiction, is comprised of a combination of equipment, procedures, rules, laws, regulations and people. As with the California Top-to-Bottom Review (TTBR), the EVEREST study unfortunately fails to consider the procedures associated with the entire election process, procedures that are integral to voting system security.
- The study does not provide a balanced approach. The EVEREST study claimed serious system flaws and failures when attacked by a "Red Team", yet failed to balance this with a "Blue Team" defense. The EVEREST Red Team attacks were carried out by computer and security experts who had unlimited access to the voting equipment and full knowledge of its design. Had it been

- implemented, a Blue Team defense would have provided balance by considering a real-world election environment, with Ohio's documented election procedures in place. In simulating an attack on any system or process via a Red Team assault, and in order to replicate "real world" conditions, it is critical that assessors also provide a Blue Team that represents the components in place designed to protect the system.
- The study is not objective. While security risks can be difficult to quantify – their likelihood, impact, etc. – the study makes no attempt to do so in a fashion that objectively compares one voting system to another or existing voting systems to their alternatives. How would a central count optical scan system, or one involving hand counted paper ballots, fare were it subjected to the same unfettered access by attackers? The EVEREST report is silent on this critical question.

The study's shortcomings can be found in the EVEREST report itself. The Executive Report of 14-Dec-2007 contains "group summary statements" from teams composed of Boards of Election (BOE) officials. In their group summary statement in response to the academic report on Premier, the BOE officials attempt to put the findings in context. Some of their comments are quoted here: (Executive Report at 45)

- "The BOE officials concluded that these sections lacked sufficient evidence relating to real-life situations in which an attacker could circumvent the security of the voting system."
- "The lack of performing these tests in real-life settings provided enough skepticism to cause the BOE officials to question the outcomes as fact-based realities."
- "Generally speaking, the BOE officials found that the report supports a certain political spectrum that believes that all electronic voting equipment is unsafe and evil."
- "The amount of mechanical errors contained within the report caused the BOE officials to question the validity of certain assertions, but it was not sufficient to compromise the credibility of the report. The study is based on clinical testing with a limited view."

It is important to note that the involvement of Ohio election administrators, while laudable, was severely constrained and limited to a review of the study that had already been completed. EVEREST would have benefited greatly from hands-on, direct involvement by these election officials in the preparation and execution of the study

from start to finish. This did not occur, and so their comments inevitably are essentially footnotes, and afforded limited weight in the final report.

The remainder of this document will address some of the key criticisms directed at the Premier voting system, as deployed in Ohio. Readers should note that an alleged vulnerability included in the EVEREST report, frequently cited by the Secretary and included in news coverage, involving an attack on the voting system using a magnet and a Personal Digital Assistant (PDA), was NOT an attack on a Premier voting system. Premier's touch screen terminals have no such vulnerability, nor was one alleged in the EVEREST study.

EVEREST Report on Premier Equipment

Note: This section focuses mainly (though not exclusively) on the EVEREST academic report, found at the following URL:
<http://www.sos.state.oh.us/sos/info/EVEREST/14-AcademicFinalEVERESTReport.pdf>

The EVEREST study reiterates evaluations done on Premier systems in past studies, one of the most recent being the California TTBR. The study also includes Premier components not previously reviewed, the Election Media Processor (EMP), ExpressPoll, and Voter Card Encoder along with a component unique to Ohio, Digital Guardian. For these and previously evaluated components the report identifies a number of new issues.

The EVEREST report has some general flaws and inaccuracies worth noting:

- The report completely fails to describe (and at times the authors seem unaware of) the procedures associated with operating the election equipment in Ohio, yet states, "Our analysis suggests that the Premier system lacks the technical protections necessary to guarantee a trustworthy election under operational conditions." (Similar statements are made for the other vendors' systems.) Ignoring procedure [operational conditions], one could fail any system, yet this sentence, given in the Executive Summary, is *the thesis* of the report. This academic study was performed by groups of computer scientists and security consultants over the course of eight weeks in an exercise that included **no physical or operational security controls**. As previously noted, one can only imagine the immense vulnerabilities that would have been identified had a

- central count optical scan system been subjected to the same conditions. A far more realistic and useful study would have applied Ohio election laws and documented security and usage procedures to the evaluation.
- The report describes serious flaws in the Premier system which are simply untrue, and was published without being vetted by the vendors. Perhaps the most glaring example of this is in the Systest Technical Report (section 3.4.1.15.3 of 17-SystestFinalTechnicalReportRedacted.pdf), which incorrectly states that the AV-TSx deletes votes without warning as its memory card reaches capacity. This incorrect conclusion was based on a misinterpretation of observations during testing. This sort of sloppiness brings the scientific credibility of the report into question. Had the authors checked their facts, they would have found that the AV-TSx never deletes votes from the current election. If memory runs low, the AV-TSx first reclaims memory consumed by previously deleted elections, then reclaims space from previous election results and data. If unable to reclaim sufficient memory for additional votes, the AV-TSx will display an out of memory message and will no longer accept votes. It is also worth noting that while a typical DRE will collect perhaps 200 votes in a twelve-hour election day, typical memory capacities allow for *tens of thousands to hundreds of thousands* of votes.
 - The report overstates or oversimplifies flaws. For example, “14.8.8 A voter can cast an unlimited number of votes without any tools or knowledge,” is a compelling sound bite for a reader skimming the report. However, this statement simply isn’t true, and is certainly not possible in the context of a polling place. The attack described requires specialized knowledge to even begin, exploiting a vulnerability that no longer exists in current AV-TSx firmware, which was not reviewed in this study. Beyond the specialized knowledge and techniques described in the report, the attacker would need a stack of voter cards to continue the attack. Even in a busy polling place, the excessive time, card reader clicking, and loud printer operation would easily alert poll workers to the activity.
 - The report makes distorted and untrue claims based on speculation. One example states, “A voter may be able to gain control of the AV-TSx.” (13.3.17). In fact, it is impossible to do so, as admitted in the technical detail: “If additional characters [...] could be entered, it would be possible to use this vulnerability to execute malicious code.” In fact, additional characters cannot be entered. The only fields seen by a voter

- are the checkbox next to each candidate name and the write-in candidate text. Accessing the field described in this attack requires first executing another attack which has already been eliminated in firmware not reviewed by the state of Ohio. While Premier concedes the code involved is not well implemented (and will correct this), the report's characterization of this and other issues is, at best, misleading.
- The report includes three issues where it states "Our evaluation has not conclusively confirmed or denied this issue. Due to time limitations and lack of access to specific un-redacted reports, we were unable to identify the vulnerabilities found in the previous studies. However, we have no reason to believe the vulnerabilities do not exist." In truth, the EVEREST reviewers had weeks to execute attacks, far more extensive and prolonged access to the voting system than one could expect any malicious attacker to enjoy. And, in fact, Premier provided the EVEREST team access to all private reports. Moreover, the California report, which previously identified these issues, contained mitigation procedures, while the EVEREST report, in some cases, did not. Given these realities, the reviewers' failure to confirm the viability of such attacks itself speaks volumes about how unrealistic such an attack scenario truly is.

Election Media Processor (EMP)

The EVEREST report includes the first published evaluation of the EMP. It enumerates eleven security issues (14.1).

The report first cites an attack on an unencrypted/unauthenticated file stored on the PCMCIA card. This vulnerability has been addressed in EMP 4.7.1 as part of the Assure 1.2 software enhancements which will soon complete federal certification, and was not reviewed by the state of Ohio.

It is correct that there is only one jurisdiction-wide data key. However, the EVEREST team appears to misunderstand how the system functions. The data keys are not associated with the machines, they are associated with the data – specifically the data stored on the PCMCIA cards. The machines themselves are interchangeable. The report's procedural mitigation strategy to "enter a unique Data Key for each card that is encoded" would in fact render the system inoperable. If the system were designed to have different

data keys for every PCMCIA card, this would imply that the EMP would need to have knowledge of every key. This would not be practical.

The process of updating keys on EMP units is conducted by trained users. Just as it is incumbent of any user of a computer system to pick a secure login password, it is equally incumbent on jurisdictions using Key Card Tool and EMP to pick strong keys. Failing to do so is not an "attack"; it is simply bad operational policy. Premier does not dictate jurisdictions' choice of keys. Users are free to choose (and are responsible for) their own password policies.

Entering a malformed IP address into EMP can indeed result in poor application behavior, but does not represent any kind of "vulnerability" as asserted by the reviewers.

The report's further observations (14.1.5 through 14.1.10), while containing incidental factual inaccuracies, are fundamentally sound. Improvements in these areas are planned for a future EMP release. EMP should only be operated by trusted individuals on a closed network.

We disagree, however, with the assertion (in 14.1.6) that "an AV-TSX can never be attached to a network." While it is true that an attacker may gain access to the client.pem file on an EMP unit, it does not follow that an AV-TSX cannot or should not be used to upload election results over a closed network. The reviewers seem to have the mistaken notion that voting results are uploaded over a "semi-public" network (one not necessarily connected to the outside world, but possibly connected to other county computers and workstations). This is absolutely not the case in Ohio, and is very strongly discouraged in Premier documentation. Instead, data is uploaded over a closed network connecting Premier voting equipment, an optional printer and *nothing else*. For larger jurisdictions, this network is typically constructed within a secured area.

GEMS Server

Given its role as the source of election definitions and the central vote counting component, the GEMS server is arguably the most critical component of Premier's election system. Premier's position is that the GEMS server itself must be protected, as should any private, corporate, or government computer containing critical information. The most effective protections are procedural and are well established

– restricting physical access, using strong passwords, operating the systems with appropriate supervision, and using the computer only for its intended purpose (i.e., not installing extra software). It is worth noting that any election management system used to tabulate and report results from centrally counted optical scan ballots, if it is to be kept secure, must be protected by the same types of physical and procedural safeguards. The burden of maintaining these procedures necessarily falls on the county Boards of Elections, who are entrusted with administering a fair election. There are number of industry standards as well as best practices and auditing methods for establishing these procedures.

Many GEMS criticisms, and many misguided measures to protect GEMS servers, arise when one questions or doubts the integrity of the election officials. (See discussion of Digital Guardian below.) That said, blind trust of election officials is not the only recourse. Maintaining proper separation of user accounts, and proper supervision during GEMS server operation, provides safeguards against corruption.

The report correctly identifies a number of flaws in the code and Premier has and will continue to improve the code quality. The report identifies interesting attacks that suggest real potential for improvement, and Premier will indeed consider these.

Premier will continue to evaluate all serious security criticisms of GEMS and will continue to improve it as a product. While some of these improvements will involve significant time and effort, we remain committed to implementing them.

AVOS

Premier’s AVOS Precinct Optical Scan counter has been effectively studied in the past. Premier has implemented improvements based on past studies, addressing the most critical vulnerabilities. Revised firmware that significantly strengthens system security is being used in other jurisdictions not requiring federal certification, and will be available for use in Ohio when the next round of federal certification is complete. This firmware was not reviewed by the state of Ohio.

From an engineering perspective the AVOS hardware is a fairly old design, with significantly less memory and computing power than more recently designed systems. This has limited the range and

number of improvements that can be made. [For example, space for audit records is indeed limited.]

For those issues which cannot be eliminated by a system change alone, jurisdictions will continue with the procedural safeguards they have successfully used over the life of the AVOS product.

For example, printing and posting the vote totals at the opening and closing of polls renders moot any attack on the memory card during transport or at the central counting location. The publicly posted precinct totals provide a simple, yet extremely effective, check against the consolidated totals reported at election headquarters. Printing the audit log at the close of polls provides a backup in case the memory card is suspected of compromise. These are basic best practices employed by OS jurisdictions throughout the nation. It is worth noting that no such integrity checking at the precinct level would be available with a centrally tabulated voting system.

The EVEREST team identifies three new (not previously reported) issues. Item (14.4.2), "The AV-OS PC ballot box collecting votes allows vote order to be reconstructed" is not true. The investigators designed a flawed test procedure, running only ten ballots, too few to establish a valid conclusion about the arrangement of the typical hundreds of ballots inserted on election day. They also fail to recognize (or report) that the ballots may fall into one of two bins depending on the system configuration, further obfuscating the order of ballot insertion. Had election officials been allowed to participate in the study as it was being performed, they could have corrected this flawed approach. Instead, BOE Officials could only comment on the procedure after the report was written, stating in the Executive Summary Report, page 45, "the Academic researchers tested voter privacy by stacking ten ballots in the ballot box. The BOE officials agreed that a proper sample for real-life application would be a test of 350 ballots."

The other new criticisms arise from the EVEREST study's failure to consider Ohio election procedure. While in use, i.e. any time the AVOS-PC and ballot box contain or are processing ballots, the system is to remain in a public area (the polling place) and is to be observed throughout its use (by poll workers). Thus the ballot box stuffing described in (14.4.1), which requires jamming and then manually stuffing ballots through the AVOS-PC, simply cannot be performed undetected in the polling place. The criticism that another manufacturer's ballot box key will unlock the Premier ballot box

(14.4.3) presents no vulnerability, given Ohio procedures. It is surprising to those who do not fully consider the system and Ohio procedures that the ballot box key itself plays a minimal security role. The key impedes casual or misguided opening and closing of the ballot box by poll workers. It is the continuous visual oversight and supervision that actually prevents attack. In Ohio, the ballot box is not typically used for storage and transfer of voted ballots.

Voter Card Encoder

The EVEREST study is the first public study to evaluate the voter card encoder (VCE). This device is based on a commodity product with a fairly open architecture, using a PIC processor and commonly available software development tools. The device has limited processing and memory capabilities. While the EVEREST study enumerates four security vulnerabilities, the team ignores handling procedures, as practiced in Ohio, and misunderstands the smart card protocols, leading to faulty conclusions.

The VCE is essentially equivalent to a stack of ballots, and must be physically protected along with other documents and materials at the polling place. Poll workers typically wear these devices on lanyards around their necks, and (if properly trained) are instructed on the importance of securing them at all times. (Similarly, in a central count OS environment it is critical that blank ballots be physically protected.) Poll workers also maintain physical control of the voter smart cards, although it is certainly possible for a voter, if not intercepted, to exit the polling place with a voter card.

Due to its open architecture, it is indeed possible to reprogram a VCE. However, the reprogrammed device will not be able to create voter cards without the proper security key. This key is not available at the polling place.

The reviewers' misunderstanding of the smart card protocol, pervasive throughout the report, is again indicated in 14.5.4, where the report laments the acceptance of the default smart card key. The report confuses the default key with the factory smart card key, and thus misses the mark in its criticism. The factory key is in effect for a card that has not yet been programmed or a card that has been erased – as in a new or previously voted voter card. Indeed, the VCE accepts this card (with the factory key) as blank and then programs it as a voter card. The VCE locks the card with the security key. The default key is

not used anywhere in this process. Thus, the criticism offered does not represent a vulnerability in the system.

ExpressPoll

The EVEREST team identifies seven issues with the ExpressPoll device. The first two issues, the second being the more serious, will be addressed in future releases of the device's operating system, just as it has been addressed for the AV-TSx.

The team identifies four issues which stem from how the voter databases are stored. This implementation places the burden of protecting voter data on the poll workers and the physical security procedures surrounding the use of the ExpressPoll. These procedures are similar to those used to protect conventional paper poll books. That said, Premier acknowledges that the use of digital media introduces new attack vectors and we are considering appropriate measures to improve the general system security for the ExpressPoll. Note that the attacks outlined by the team require physical access to the device, so well designed and implemented procedures can sufficiently protect this high-value device.

Finally, the team identifies (14.6.5) a vulnerability during transit in which an attacker applies a high voltage charge to the unit through its charging port. Although the report documents the attack well, it ignores the fact that these types of attacks can be executed on most electronic equipment and without physical contact. Indeed, if afforded uncontrolled access, an attacker can destroy a unit with a hammer, or destroy or steal a paper poll book. All voting system components are high-value, and must be protected in transit. A typical strategy is to maintain a documented chain of custody.

Digital Guardian

Digital Guardian is a software component used on the GEMS servers as an added security measure. Many of the GEMS server exploits considered by EVEREST require that Digital Guardian be disabled. The unfettered Red Team found ways to do this, but these required knowledge of the BIOS password or access to the interior of the GEMS server.

The justification for using Digital Guardian in Ohio is fairly complex. The previous Ohio SOS administration commissioned studies of potential voting systems for Ohio, including Premier's. The Compuware study identified concerns with the GEMS server using a database that could be accessed via standard Microsoft Office tools, the Jet database. Premier's position was, and continues to be, that standard methods of protecting any computer – restricting physical access, using strong passwords, maintaining separate user accounts, would prevent unauthorized use of the GEMS server. Our security model depended on a certain level of trust of elections personnel and their procedures to protect the GEMS server.

Right or wrong, trust in election officials is no longer commonly accepted. At the request of the previous Secretary of State's administration, Premier identified Digital Guardian as a product that could be configured to fine-tune restrictions and access on the system. It was not ideally suited to Ohio's county voting systems, as they are not connected to a central management server. (Please note, Digital Guardian is an excellent and successful product in many areas.) But in effect, it gave the Ohio SOS the "keys" to all (Premier) voting systems in Ohio, and this satisfied Compuware, and therefore the Ohio SOS.

Despite the cumbersome and costly (for both Premier and the State of Ohio) management of this system, it remains in place. No other Premier customer in any other jurisdiction has ever used this product to protect their GEMS servers.

Premier's position on Digital Guardian is that there are other more effective, sensible, and less burdensome methods to protect the GEMS servers in Ohio. These should be seriously considered, with realistic understanding of threats.

Nevertheless, with a properly protected BIOS password and reasonable industry standard controls on physical access to the GEMS server, Digital Guardian cannot be "taken down". Thus, the numerous GEMS server attacks identified by EVEREST that depend on installing and running various intrusion packages are not valid.

AVTSx

Perhaps the most serious criticism of the AVTSx is the capability to update its firmware without proper credentials (13.3.1, 13.3.2). While Premier recognizes this as a vulnerability, and has eliminated it in our newest version of firmware, we do object to the EVEREST report's oversimplification of using this mechanism to attack the entire voting system. Without demonstrating or even properly describing the structure of malicious firmware, the EVEREST report hypothesizes, but could not confirm through testing, viral propagation of malware from the voting machines up through the EMP and into GEMS. Although Premier takes seriously and seeks to correct any potential vulnerability in its systems, we believe a "bright line" must be drawn between attacks that can be successfully executed in testing and those that are solely theoretical, and have not been accomplished even in a laboratory setting of unfettered access. Nonetheless, our current firmware, now in federal certification and not considered in the EVEREST study, eliminates the possibility of unauthorized updates.

The report claims – fantastically – that "a voter can cast an unlimited number of votes without special tools or knowledge" (14.8.8). Yet the tools, knowledge, and access required are extraordinary. (We fully discuss this item in this response's Introduction section, where we discuss the EVEREST report's general flaws and inaccuracies.)

There are numerous observations related to the code's data integrity checking, buffer overflow vulnerabilities, and general defectiveness leading to crashes or injections of viruses. Yet the authors have been unable to demonstrate a viable attack. Many of these areas have been reviewed and improved in the years since the examined code base was introduced. This is an ongoing process, and much has been improved in the current version being considered for federal certification.

The report criticizes Premier for the use of a "default key", chiding that there is no such thing as a default secret. We agree. Premier documentation strongly urges election officials to generate their own election-specific keys (as is done in Ohio) and provides a clear visual indicator (a broken key) on the voting machine in case the default key is still in place.

The use of a single key over multiple precincts, or over an entire county is an important operational question for election administrators.

While multiple keys arguably provide higher security, their use adds administrative burden, which may lead to sloppy handling of keys by multiple parties or loss of service if the keying procedures are botched. The Premier system does not dictate that all machines have the same key, and, in fact, provides the capability to create multitudes of specific keys.

The report's treatment of the VVPAT system identifies or reiterates a number of concerns, all of which involve physical attack on the unit. These attacks are particularly problematic in Ohio, which treats the paper audit trail as the official ballot in case of disputes. Premier will consider these physical attacks in planning improvements to this product. However, we note that it is a matter of opinion as to how difficult these would be to carry out undetected. The report characterizes the attacks as "easy". In the real-world context of a polling place, Premier disagrees.

The report ignores a key mitigation strategy in case of the destruction of an audit roll: The audit trail may be reconstructed by having the machine print a new copy. Understandably, this mitigation is of little use for those who mistrust the electronic records. However, up to the point of destruction, each and every voter would have had a chance to verify that their vote was properly recorded at the time and to alert a poll worker if not – one of the tenets of using a VVPAT.

Conclusion

Premier fully cooperated with the Everest reviewers and we welcome objective evaluation of Ohio's election system. We respect the ingenuity and hard work of the EVEREST team's investigators. They have provided us with useful insights. Our engineering group has reviewed and will continue to consider mitigations for every potential vulnerability identified in the study. We have supported the EVEREST study in hopes of seeing useful results to help improve all voting systems and we will continue to provide support to future studies.

Yet, unfortunately, the EVEREST study failed to meet its publicly stated goals by excluding consideration of real-world procedures in Ohio, procedures mandated and managed by the Ohio SOS and state law.

Any public voting system – paper, punch card, optical scan, lever, electronic – can and will be found vulnerable to attack if left open to full access in an environment that ignores physical and procedural

security. This is certainly the case with the central count optical scan model that the Secretary of State has advocated and embraced. There is no voting system in existence, nor is it possible to envision one, that could withstand unfettered and unmonitored attacks in a laboratory setting whereby dozens of computer experts spend hundreds of hours, with full access, attempting to infiltrate the system.

Premier acknowledges there are authentic potential vulnerabilities, and will continue to investigate and address each one that is valid. Unfortunately, any system, including ours, will never be free of vulnerabilities, especially without proper procedures in place. Also problematic is the time required to certify system improvements: The EVEREST authors puzzle over Premier's apparent lack of improvement in light of one study after another, stating the only "reasonable" conclusion is failed engineering approaches. (Academic Report, Chapter 10) Yet, they fail to comprehend that the current Ohio-certified versions date back to early 2004. We have continually improved our systems, and these are in certification now.

The report's executive summary gives the study's objectives, beginning with this paragraph on page 6:

The ultimate objective of Project EVEREST is to improve the integrity of Ohio elections for federal office, and state and local offices and issues, and provide the citizenry with increased confidence and trust in our elections system.

Unless the reader approaches this study's reports with healthy skepticism, objectivity, intellectual honesty, and practical knowledge of how counties run elections in Ohio, the study may yield just the opposite: Election integrity could actually decrease as policy decisions are made based on this flawed study. Confidence in current voting systems may be undermined as voters are presented with claims of potential threats, many of them unsubstantiated, without any scientific analysis of whether alternative systems would create new, and easier to exploit, vulnerabilities.

Based on the EVEREST study, the Ohio Secretary of State has recommended central optical scan of paper ballots, and is considering combined vote centers which would be available for weeks prior to an election. Ironically, the new avenues of attack against centralized stacks of paper stored for weeks and transported daily within Ohio's 88 counties are unquestionably much less sophisticated, and therefore much more achievable by a wider range of attackers, than those

against electronic equipment. These new avenues of attack have not been studied or properly quantified.