BerkeleyLaw
UNIVERSITY OF CALIFORNIA

University of California, Berkel
School of Law
396 Simon Hall
Tel 510-410-6964
Fax 510-643-4625

April 28, 2008

Mr. Matt Masterson
U.S. Election Assistance Commission
1225 New York Avenue, Suite 1100
Washington, DC 20005

Dear Mr. Masterson:

Thank you for the opportunity to comment on the Election Assistance Commission's proposed Statement of Work concerning voting system risk assessment. We are pleased to submit the attached comments in response to the request for information. We would be happy to answer any questions that you or Commission members or staff have about our comments. Please notify me if we may be of further assistance.

Sincerely yours,

/s/
Aaron J. Burstein
TRUST and ACCURATE Research Fellow
UC Berkeley School of Law

/s/
Joseph Lorenzo Hall
PhD Candidate
School of Information
UC Berkeley

/s/
Deirdre K. Mulligan
Clinical Professor of Law
Director, Samuelson Law, Technology & Public Policy Clinic
UC Berkeley School of Law

/s/ David A. Wagner
Associate Professor
Computer Science Division
UC Berkeley

Encl: Comment on the U.S. Election Assistance Commission's Request for Information Regarding Voting Systems Risk Assessment Support

Comment on the U.S. Election Assistance Commission's Request for Information Regarding Voting Systems Risk Assessment Support

April 28, 2008

We appreciate the opportunity to comment on the Election Assistance Commission's Request for Information (RFI) concerning the statement of work (SOW) for a voting systems risk assessment. A structured risk assessment model for voting systems is badly needed; a model that is as complete and realistic as possible would substantially benefit the development of voting systems guidelines, voting system technology, and election administration. Situating an analysis of threats to voting systems within the context of election process models is an ambitious project, and we are encouraged to see the Commission providing the time and funding for it.

We write to suggest modifications to the SOW that we believe would improve the EAC's effort substantively and procedurally. These suggestions fall into five categories: (1) recognizing that a voting system risk assessment is likely to be incomplete; (2) explaining how the EAC will use the quantitative versus qualitative parts of the assessment; (3) recognizing the trade-offs involved in requiring models developed during the risk assessment to be fully and independently usable by non-experts; (4) affording public access to risk assessment tools; and (5) eliminating conflicts of interest among members of the risk assessment team.

At the outset, we offer a general framework developed in the field of risk analysis, in order to place the SOW in broader perspective and to establish consistent terminology. This framework identifies three stages of risk analysis:[1]

1. Hazard identification (What can go wrong? What are the impacts?)

2. Risk measurement (What are the probabilities?)

3. Risk management (What are the mitigations?)

[1] *See* Alfredo Garcia & Barry Horowitz, *The Potential for Underinvestment in Internet Security: Implications for Regulatory Policy*, 31 J. REG. ECON. 37, 50 n.4 (2006); Paul B. Thompson, *Ecological Risks of Transgenic Plants: A Framework for Assessment and Conceptual Issues*, 16, 18-19 *in Issues in Environmental Science and Technology No. 21: Sustainability in Agriculture* (Royal Society of Chemistry) (2005).

Recognize the Incompleteness of Hazard Identification. Generally speaking, the usefulness of the EAC's proposed risk assessment depends heavily upon the completeness of hazard identification, that is, how completely it identifies voting system threats and vulnerabilities. We urge the EAC to recognize the difficulty of achieving completeness along either of these dimensions.

The model of the Federal Information Security Management Act (FISMA) framework—which we do not interpret to require the proposed risk assessment—is consistent with the general framework outlined above; but it does not guarantee that hazard identification will be complete. The FISMA framework involves identifying specific information security threats, stating the likelihood of each, implementing cost-effective policies and procedures to reduce the risk posed by these threats, and periodically testing and evaluating these measures.[2]

But neither FISMA nor the supporting NIST documentation provide a general means of assessing the completeness of hazard identification. Instead, the framework supports choosing security mechanisms (management, administrative, technical, and physical standards and guidelines) that are appropriate based on the hazards that are identified, and defining agency compliance with a plan to implement those mechanisms. In practice, this framework has led to a heavy focus on compliance rather than the actual security performance of the agency.[3] This is a useful approach in administering an information system that faces well-understood security threats—for example, by ensuring that a particular information system installation is properly configured and patched.

But its value is less clear for abstract models of voting systems. A great deal remains to be learned about threats—both technical and non-technical—to these systems. Consequently, the completeness of any voting system hazard assessment is a major issue, and any use a compliance-based approach should take this into account. Experience with existing voting system guidelines and actual voting systems has shown that emphasizing compliance to the exclusion of other security evaluation approaches can leave serious vulnerabilities undetected for long periods of time.[4] Moreover, this experience has shown that it is difficult to identify hazards to a voting system without meaningful access to that system (e.g., access to equipment, source code, documentation, and knowledge of the procedures and policies governing its use in specific contexts). It is unclear whether

[2] 44 U.S.C. § 3544(a)(2).

[3] *See* James A. Lewis, Testimony before the House Committee on Oversight and Government Reform Subcommittee on Government Management, Organization, and Procurement and the Subcommittee on Information Policy, Census, and National Archives testimony, June 7, 2007 ("FISMA is a direct measurement of compliance with processes and an indirect measure of performance.").

[4] Florida State University Security and Assurance in Information Technology Laboratory, *Software Reviews and Security Analyses of Florida Voting Systems*, February 2008, *at* `http://www.sait.fsu.edu/research/evoting/index.shtml`; Patrick McDaniel et al., *EVEREST: Evaluation and Validation of Election-Related Equipment, Standards and Testing* (Academic Final Report), December 2007, *at* `http://www.sos.state.oh.us/sos/info/EVEREST/14-AcademicFinalEVERESTReport.pdf`; California Secretary of State, *Top-To-Bottom Review of California's Voting Systems*, August 2007, *at* `http://www.sos.ca.gov/elections/elections_vsr.htm`; Ariel J. Feldman, J. Alex Halderman & Edward W. Felten, *Security Analysis of the Diebold AccuVote-TS Voting Machine*, *in Proceedings of USENIX/ACCURATE Electronic Voting Technology Workshop*, August 2007, *at* `http://www.usenix.org/events/evt07/tech/full_papers/feldman/feldman.pdf`.

the Contractor will be given access to actual voting equipment under the SOW; the EAC should add a section to Phase I of the SOW to clarify that studying actual voting equipment is an element of creating voting system reference models.

The difficulties of hazard identification are even more acute for *future* voting systems. Any assessment of systems that do not yet exist will likely be either speculative, conjectural, and incomplete (if one tries to examine the risks in detail) or so broad that it is ineffective for drawing conclusions and difficult to estimate quantitatively the risks (if one tries to look at broad categories).

The EAC should recognize these limitations when it uses the risk assessment to determine the VVSG's security requirements. The EAC might regard the hazard identification as a baseline or lower bound of threats that all voting systems should defend against. The VVSG draft, of course, contains mechanisms for identifying and addressing threats that ex ante hazard identification might miss. Specifically, the software independence requirement, the adversarial vulnerability testing,[5] and the volume testing requirements would provide proven means of discovering vulnerabilities prior to certification as well as a general approach to prevent *undetected* errors from changing the outcome of an election.[6]

We recommend the following changes to the SOW to help increase the completeness of the risk assessment's hazard identification.

- Require public review of the hazard identitifcation to bring omitted threats into the risk assessment. This review could be achieved by requiring periodic publication of working drafts of the hazard identification, followed by a public comment period, or through public scientific workshops held to discuss the model. As an example of the workshop model, we point to NIST's development of the Advanced Encryption Standard (AES), which included a series of workshops to facilitate public, scientific discussion of the standard as it developed.[7]

- Require that the Contractor include hazards that are identified by others (e.g., through academic research or state-level testing) while Phases II and III are underway. Though we recognize the need to balance completeness with finite time and resources for the assessment, the current SOW draft leans too far away from updating the threat model during the term of the contract.

Clarify the Role for Risk Measurement. The prospect of quantitatively assessing risk to voting systems is attractive,[8] but we believe this task must be approached with caution. To clarify what we mean by risk measurement (the second element of our risk

[5] The VVSG draft refers to this form of testing as "open-ended vulnerability testing," a term that, unfortunately, suggests to some that this form of testing is standardless and impossible to satisfy.

[6] These issues are explored in greater depth in ACCURATE's forthcoming comments on the VVSG draft.

[7] *See* NIST, *Advanced Encryption Standard*, *at* `http://csrc.nist.gov/publications/nistbul/itl97-02.txt`.

[8] The core tasks of developing and refining risk assessments require the production of quantitative assessments of risk and the use of mathematical modeling (SOW §§ 4.12 and 4.13).

analysis framework),[9] we offer the following algorithm. Measuring risk for a voting system would involve three steps:

1. Identifying all possible failure modes of the system, i.e., conducting the hazard identification discussed above.

2. Computing damage from each failure mode, weighted by the probability ifs occurrence.

3. Summing (2) over all failure modes.

Steps (1) and (2) of this approach counsel taking a cautious approach to voting system risk measurement. As discussed above, any hazard identification (Step (1)) is likely to be incomplete, which can underestimate risk by omitting threats.

Furthermore, Step (2) involves quantities that are difficult to estimate. To take the probability-of-attack element first, the computation would need to account for the motivations of attackers, including the possibility that an insider will leave backdoors to make systems easier to attack in the field.[10] Also, successful attacks might not be randomly distributed—attackers might focus on a few jurisdictions—making the total damage from a particular failure mode hard to estimate. In addition, viral attacks against voting systems could radically reduce the cost of corrupting a large number of votes or voting systems.[11] We are not aware of any sources of data that would provide a basis for meaningful estimates of these quantities. On the matter of damages, it is unclear how to put a price on an accident or successful attack against a voting system that calls into question the results of an election.[12] The full impact of a significant error or successful attack may become manifest as a loss of voter confidence and thus be difficult to measure precisely.[13]

Still, the risk assessment should tie the hazard identification to a sense of what it would take to perpetrate a successful attack, or what it would take for errors in a voting system to change the result of an election. A thorough hazard identification, combined with reasoning about the motivations of attackers, the impacts of accidents or successful attacks, and other factors—some of which may be qualitative—can lead to helpful threat prioritization and recommendations for mitigations.[14] A very helpful step in this direction

[9] We will use "risk quantification" interchangeably with "risk measurement."

[10] *See* Douglas W. Jones, *Threats to Voting Systems*, Position Paper for the NIST Workshop on Developing an Analysis of Threats to Voting Systems, Gaithersburg, MD, Oct. 7, 2005, *available at* `http://www.cs.uiowa.edu/~jones/voting/nist2005.shtml` (noting that "[t]he fact that so many of the costs are fuzzy poses a serious problem" to quantitative evaluation of risk).

[11] Examples of such attacks are discussed in the California Top-to-Bottom reports and by Feldman, Halderman, and Felten, cited in note 4.

[12] Accidents and errors are part of the threat landscape and, presumably, must be subject to quantitative analysis. *See* Jones, *supra* note 10 (noting that the use of "threat" within computer security includes both malicious acts and accidents or mistakes.

[13] For more on this point, see Lawrence D. Norden, Written Comments Submitted Before U.S. Election Assistance Commission Voting Advocate Roundtable Discussion, April 24, 2008.

[14] The principal example in the context of voting systems is Lawrence Norden et al., THE MACHINERY OF DEMOCRACY: PROTECTING ELECTIONS IN AN ELECTRONIC WORLD (Brennan Center for

was taken in the Brennan Center's *Machinery of Democracy*, which used "attack team size"—the number of informed participants necessary to carry out an attack—to measure the difficulty of an attack.[15] The eventual Request for Proposals should cite this report explicitly as a foundation for the Contractor's work. A risk assessment framework built on such a metric would distinguish difficult-to-conceal attacks (such as those that require the cooperation of poll workers in many precincts) from those that are far more easily hidden (such as malicious code written by a single rogue programmer).

We recommend that this framework be expanded to include at least a rough measure of what would have to be done to avoid noticing *accidental* errors in a voting system. As discussed above, both accidents and attacks are relevant voting system hazards.

In summary, we recommend changing the SOW's focus from risk measurement guided by the algorithm sketched above, with its emphasis on the likelihood of a voting system failing in specific ways, to a focus on the following questions:

- What can go wrong in a voting system?
- What would it take to make something go wrong (intentionally) or to avoid noticing that something had gone wrong (through attack or accident)?
- What can be done to prevent such events from occurring, to detect them if they do, and to mitigate the damage?

Reconsider a One-Size-Fits-All Assessment Tool. Our third area of concern is that the SOW seems to set an impossible goal by requiring that the tools developed during the risk assessment be usable by the EAC and all other stakeholders "without the assistance of specialized experts" (§ 1.0; § 4.15). Requiring the use of models that non-experts can use independently might unduly constrain the kinds of analysis that the Contractor may perform, and thus make the risk assessment less informative than it could be. This difficulty is exacerbated by the requirement of a one-size-fits-all product; the models that the EAC could use independently, for example, might be more sophisticated than those that a local election official could use. The EAC should consider revising the SOW to specify a final set of tools that would be appropriate for specific stakeholders, e.g., local election officials, voting system manufacturers, and the EAC itself.

Make all risk assessment tools, with adequate documentation, publicly available. Aside from specifying that the risk assessment must be usable by the EAC and election officials, and requiring the Contractor to brief EAC and NIST,[16] the SOW does not state how the product will be distributed. Making the models that result from this

Justice ed., 2006), *available at* http://www.brennancenter.org/content/resource/machinery_of_democracy_protecting_elections_in_an_electronic_world/. A similar mode of analysis is common in the computer security community for architectural risk analysis.

[15] *See* Norden et al., Machinery of Democracy, *supra* note 14, at 24-25.

[16] See our suggestions above to expand peer review to include public comments or public workshops on the hazard identification element of the risk assessment.

effort public would advance the public interest in voting system security (as well as accuracy, reliability, etc.) by allowing the widest possible scrutiny and use. If the models are implemented in computer programs, the source code should be made publicly available for use and further development. Similarly, the EAC should require thorough documentation of the code and make the documentation publicly available. These steps would not only facilitate further work on voting system risk assessment but also enhance the transparency of the EAC's decisions on the VVSG's security requirements.

Prohibit Conflicts of Interest. Finally, the SOW should supplement its requirement of a "well-qualified and broadly-based team" with a prohibition on financial or personal conflicts of interest among team members. It is crucial to the integrity of this effort that neither the individuals on the team nor the entity (or entities) that are parties to the contract have current or recent ties to voting system manufacturers or test labs. One possibility for avoiding these conflicts entirely is to have NIST perform or take an active role in the risk assessment. We recommend that the EAC explore this possibility and state in the final request for proposals whether NIST may become involved in the assessment.