

End-to-end crypto works. Now what?

Dan S. Wallach, Rice University

1 Introduction

The disconnect between the research community and the practice of voting is stark. We now have a variety of techniques that can provably catch tabulation errors, verifying that votes are recorded as they were cast and tabulated as they were recorded. The mathematics are sound. Prototype systems have been implemented [10, 5] and in some cases field tested [1]. Despite this, none of the major vendors of voting technologies have made any moves toward adopting e2e techniques in their systems. None of the standards organizations or regulatory bodies have moved toward requiring these technologies. The best we have is the EAC’s “innovation class” that at least *allows* these things. This paper attempts to address the question of what’s missing. There are several possible reasons that we might consider.

Comprehension. A common lament is whether e2e techniques can be understood by the non-cryptographer. What good is a technology that requires a doctorate in a specialized field to understand? Even the various paper-based “non-cryptographic” schemes like Prêt à Voter [4] or ThreeBallot [9] rely on fairly subtle properties to guarantee their fairness against coercion attacks [8].

Privacy. E2e techniques protect the *integrity* of votes against tampering, but have weaker guarantees for voters’ *privacy* [7, 11]. Research to address these problems (see, e.g., [6]) is only just beginning.

Discredited DREs play a significant role as well. The message the current-generation direct recording electronic (DRE) voting systems are wildly insecure has finally taken root. Anybody who denies the poor quality of their software engineering and the preponderance of powerful attacks against them, as made clear in California’s 2007 “Top to Bottom Report” and the follow-on “EVEREST” work from Ohio, is at best willfully ignoring the truth.¹ Unfortunately, this has created a popular perception that DRE system can *never* be secure, despite their potential to be otherwise.

Economic forces ultimately drive what vendors build. Where’s the demand? In the absence of the funds made available from HAVA in 2002, compounded with recessionary pressures on local budgets, local election officials are cash-strapped and largely cannot afford to buy expensive new equipment. This places the industry into a maintenance mode, since what revenue they’re getting is coming from existing support contracts on fielded equipment, rather than new procurements. To the extent that any new equipment is being adopted, the political pressure is to adopt optical scan systems, both for their lower cost and their improved security properties. E2e techniques are not on the radar.

Certification has also become a nightmare for vendors with the transition from the previous NASED process to the present-day EAC process. Demonstrably, the original NASED process was a very coarse filter. Grossly deficient equipment was nonetheless “certified” and widely deployed across the U.S. It’s great that the EAC is trying to raise the bar, but the customer demand, such as it is, appears to be not for equipment meeting the new standards but rather for equipment that’s available today, grandfathered in from the old

¹Or may even genuinely believe DREs are secure, or that they can detect tampering, despite the evidence to the contrary.

standards. Vendors today are still shipping software from years ago because they cannot (or choose not) to meet the newer standards.

2 Breaking the Logjam

One attractive direction, taken by the Scantegrity project [3] as well as a proposal by Benaloh [2], adapt e2e techniques to existing paper ballots. Voters have a similar experience to traditional precinct-based optical scanned ballots, but they may opt to do additional work (with Scantegrity, writing down codes revealed by the invisible ink pen for later verification; with Benaloh's scheme, challenging the optical scanner to prove it performed the cryptography properly).

One challenge for the e2e community is to prove that our techniques are *usable*. There are two broad schemes for measuring this: mock elections and controlled laboratory tests. The former shows that a technique can scale and may turn up low-frequency yet significant problems. The latter can consider many experimental variations on how a particular voting scheme might work and then quantify whether there are significant variations in usability across the population of voters. Both of these techniques must be applied before any e2e technique is to get out of the lab and into production. (Arguably, this sort of work should be required for any commercial voting system as part of the certification process.) We still have the burden of demonstrating that e2e techniques will work and be acceptable in the hands of the everyday voter and poll worker.

A more important challenge is creating the economic and regulatory climate that will forcibly move vendors and municipalities away from older, flawed technologies to newer, verifiable technologies. Without another massive injection of government money, this won't be easy.

A radical but entirely feasible solution would be to restructure the U.S. voting system industry. Since vendors are unable or unwilling to produce good systems, the government can step in and create a public/private partnership. Much of the complexity in meeting the EAC's current standards comes from requirements on the software. The hardware requirements, for contrast, are relatively straightforward. This leads to a natural division of labor. A government-funded non-profit software development effort could design and ship the software for DRE and optical scan systems, as well as the back-end tabulation machinery. Vendors could then design hardware that meets the government's specifications, without having to pay for the expensive and difficult process of redesigning their legacy software to meet more stringent standards. A very similar business model has been incredibly successful for personal computers; Microsoft provides the software while other vendors ship hardware, sometimes with modest software customizations. As a nice side-benefit, standardized software would commoditize the hardware, making it easier to mix-and-match equipment and driving prices down.

Of course, vendors will resist. They will discuss how this undercuts their ability to innovate (but yet they're not really innovating). They will complain of their inability to meet customer needs and customizations (which they could certainly do under an open-source model, although their changes would need to pass muster and would then be available to other hardware vendors). They will decry the cost associated with junking perfectly good, fielded equipment (which is actually reaching the end of its service lifetime, and new equipment would be radically cheaper).² There will be inevitable allegations of patent infringement (possibly requiring the government to indemnify hardware vendors and fight these battles on their behalf).

All of these issues can and must be addressed if we truly want to see better voting technologies deployed in the U.S. If we fail to do this, we may instead see the end of the polling place as we know it, with the unfortunate and growing adoption of vote-by-mail or even vote-by-Internet schemes, which have no meaningful resistance to voter coercion. Polling-place voting is too valuable for security to surrender so easily.

²Current DRE voting system costs thousands of dollars a piece, plus maintenance fees. The hardware in a voting system needs not be significantly different from the hardware found in \$400 "netbook" computers, save having more rugged cases. Significant cost savings are eminently feasible.

References

- [1] ADIDA, B. Helios: Web-based open-audit voting. In *17th USENIX Security Symposium* (San Jose, CA, July 2008).
- [2] BENALOH, J. Administrative and public verifiability: Can we have both? In *Proceedings of the 3rd USENIX/ACCURATE Electronic Voting Technology Workshop (EVT'08)* (San Jose, CA, Aug. 2008).
- [3] CHAUM, D., CARBACK, R., CLARK, J., ESSEX, A., POPOVENIUC, S., RIVEST, R. L., RYAN, P. Y., SHEN, E., AND SHERMAN, A. T. Scantegrity II: End-to-end verifiability for optical scan election systems using invisible ink confirmation codes. In *Electronic Voting Technology Workshop 2008* (San Jose, CA, Aug. 2008).
- [4] CHAUM, D., RYAN, P. Y. A., AND SCHNEIDER, S. A. A practical, voter-verifiable election scheme. In *ESORICS '05* (Milan, Italy, 2005), pp. 118–139.
- [5] CLARKSON, M. R., CHONG, S., AND MYERS, A. C. Civitas: A secure voting system. In *IEEE Symposium on Security and Privacy* (Oakland, CA, May 2008).
- [6] FELDMAN, A. J., AND BENALOH, J. On subliminal channels in encrypt-on-cast voting systems. In *Electronic Voting Technology/Workshop on Trustworthy Elections 2009* (Montreal, Canada, Aug. 2009).
- [7] KARLOF, C., SASTRY, N., AND WAGNER, D. Cryptographic voting protocols: A systems perspective. In *USENIX Security Symposium* (Aug. 2005).
- [8] KELSEY, J., REGENSCHIED, A., MORAN, T., AND CHAUM, D. Hacking paper: Some random attacks on paper-based E2E systems. Presentation in Seminar 07311: Frontiers of Electronic Voting, 29.07.07–03.08.07, organized in The International Conference and Research Center for Computer Science (IBFI, Schloss Dagstuhl, Germany), Aug. 2007. <http://kathrin.dagstuhl.de/files/Materials/07/07311/07311.KelseyJohn.Slides.pdf>.
- [9] RIVEST, R. L., AND SMITH, W. D. Three voting protocols: ThreeBallot, VAV, and Twin. In *Proceedings of the 2nd USENIX/ACCURATE Electronic Voting Technology Workshop (EVT'07)* (Boston, MA, Aug. 2007).
- [10] SANDLER, D. R., DERR, K., AND WALLACH, D. S. VoteBox: a tamper-evident, verifiable electronic voting system. In *Proceedings of the 17th USENIX Security Symposium (USENIX Security 2008)* (San Jose, CA, 2008).
- [11] WALLACH, D. S. Voting system risk assessment via computational complexity analysis. *William & Mary Bill of Rights Journal* 17 (Dec. 2008).