February 7, 2011

Representative Joyce Peppin
Government Operations and Election Committee
503 State Office Building
100 Rev. Dr. Martin Luther King Jr. Blvd.
Saint Paul, MN 55155

Dear Ms. Peppin and members of the committee:

We are concerned about the electronic roster provisions of H.F. 210DE1 (hereafter referred to as H.F. 210). Each of us has professional expertise regarding computerized information systems; additionally, we are affiliated with Citizens for Election Integrity Minnesota (CEIMN), a nonpartisan organization that advocates for verifiable, transparent and accurate elections.

First, and most fundamentally, any electronic roster system that includes online access from precinct polling places to centralized servers is an inappropriate application of technology. Such an online electronic roster, as proposed in H.F. 210, will incur substantial costs and will risk serious disruption to the smooth operation of elections, neither of which is justified by the marginal improvements to election integrity that might be achieved. Second, the detailed provisions of H.F. 210 governing the electronic roster show numerous specific signs of inadequate vetting by knowledgeable professionals and as such would be unwise to codify. So as not to burden this letter with technical details, we have separated our more detailed comments into an addendum.

Our concerns regarding the electronic roster provisions of H.F. 210 are sufficiently grave that we would caution even those members who support other aspects of the bill against voting to move the bill forward. These provisions may seem like minor technical details compared with such hot-button issues as photo identification. However, if enacted without adequate scrutiny, they will come back to haunt Minnesotans.

During the election, each polling place currently operates independently, with centralized information only flowing in before the election (when the roster is printed) and only flowing out after the election (when computer data entry occurs).

H.F. 210 would make not one, but two changes: it would replace the paper roster with a computerized one, and it would replace autonomous operation of the polling places with continuous online connection of the polling-place computers to centralized servers.

We must emphasize that this bundling together of two changes is not inevitable: the state could move to computerized rosters while preserving autonomous operation during the election. However, H.F. 210 does not provide for such offline operation; it requires continuously online operation (other than in exceptional circumstances) with "all voting information processed by any computer in a precinct … immediately accessible to all other computers at all other connected precincts in the state" (Article 3, Section 3). This is the form of electronic roster to which we write in opposition.

Thousands of locations across the state serve as polling places on a temporary basis, with staffing provided largely by dedicated volunteers. Even if the locations and staff were permanent, providing secure, reliable communication from computers at each of these sites to centralized servers would be a major technical challenge. Any vendor who undertook that systems-integration challenge would quite rightly charge fees well in excess of the hardware costs. The temporary nature of polling places greatly exacerbates

the challenges. Our experience teaches us that distributed computer systems have their highest failure rate when first set up. We have come to expect unexpected delays in getting systems to work. At a polling place, such delays would be highly problematic and could result in frustrated voters leaving the polling place before they vote.

Nor is it much comfort that election judges, faced with technical problems, could eventually give up on getting the computers to work and fall back on the paper-roster contingency plan provided for by the bill. Deciding when to give up on troubleshooting is never easy. Moreover, this contingency plan undermines the proponents' claim that the cost of a new system would be offset by eliminating printing costs for rosters.

Any online electronic roster would inevitably prove both costly and likely to result in disruptions to the smooth conduct of elections. As such, an online electronic roster would be an inappropriate application of technology unless it achieved major advantages over simpler offline (or perhaps even paper-based) alternatives. Thus, it is worth considering the two ways in which online access might further election integrity goals.

The first advantage to an online system is that it would ensure that no voter identity was used to vote in two different precincts. This is the apparent motivation for the previously quoted requirement that information from each precinct become immediately available in all other precincts. We are not talking about one individual going from precinct to precinct, using a different assumed identity in each one; such use of assumed identities is targeted by the photo ID requirement. Instead, we are talking about someone using his or her actual identity in more than one precinct – double voting. We have found no evidence to suggest that this form of fraud is more than a hypothetical problem. Moreover, even an offline system would allow this peculiar form of fraud to be deterred through a credible threat of prosecution. The benefit here is a marginal gain in an unlikely, hypothetical situation that would in any case be deterred. We are convinced this benefit does not merit significant costs and the risk of election-day chaos.

The second advantage to an online system is that for election-day registrants, it would allow their eligibility to be verified by checking centralized databases. Unlike the case of someone voting in multiple precincts, this situation is not purely hypothetical: ineligible individuals do on rare occasion attempt to register and vote. Often, this is not intentional fraud, but rather the result of misunderstanding and might be addressed by providing clearer information. Although any amount of fraudulent registrations is bad, stamping out the tiny amount that the county attorneys have been able to identify is not worth risking the smooth operation of our election system for the rest of us.

In summary, the online election roster specified by H.F. 210 is an unwise idea at a fundamental level. Moreover, as described in the addendum to this letter, the detailed specifications contained in the bill will serve to further make life difficult for the election officials charged with implementing the system. They will need to make the best of specifications that though quite prescriptive are ambiguous, contradictory, and inadequately vetted. This struggle will make an inherently difficult and expensive task even more so. Those detailed issues need not be considered, however, because *any* online election roster system would be a significant, unjustified expense for the taxpayers and a frightening risk to the smooth conduct of our elections.

We thank you for your attention to our concerns. The first signatory will attend Tuesday's committee meeting and would be happy to answer any questions.

Sincerely,

Max Hailperin, PhD
Professor of Computer Science, Gustavus Adolphus College
Volunteer, Citizens for Election Integrity Minnesota

Stan Hilliard
Retired Specialist of 3M's IT Statistical Consulting Dept. and Co-Owner of H & H Servicco Corp
Member, CEIMN Organizing Committee

Addendum to CEIMN letter regarding H.F. 210's electronic roster provisions

There are additional trouble areas within the specific language of Article 3. The first sign that Article 3 was not adequately vetted is that two portions do not mesh properly. Section 3 would add a new section to the Minnesota Statutes, 201.225, providing standards for electronic rosters. Sections 8-15 would add an entire new chapter to the Minnesota Statues, 206A, that covers much of the same ground. The language for 206A is not simply more detailed than that for 201.225; at times, they are outright conflicting.

Take, for example, the question of how broadly centralized the electronic roster system is to be. We already quoted a passage from Section 3 that requires updates made in any precinct to be "immediately accessible to all other computers at all other connected precincts in the *state*." (Emphasis added.) Yet Section 9 defines the term "electronic roster" as a list "which shall be processed by a computer at a precinct to be immediately accessible to all other computers at all precincts in the *county*." (Emphasis added.) So which is it: the state or the county?

The confusion exists not only between sections. Even within Section 9, the definition of "election roster" with a countywide scope for information sharing is followed by a definition for "teleprocessing lines" as being "between precincts and a centralized computerized roster maintained by the secretary of state."

This confusion regarding countywide versus statewide scope is quite significant for the process by which the election roster system would be developed. Sections 8-15 assign the responsibility to the county and municipal election officials, with the secretary of state's role limited to approving the plans submitted by the local officials. Even if the scope of online information sharing were not intended to be statewide, placing so much responsibility on local officials seems like an unreasonable burden. Our experience with local election officials gives us great respect for their ability to work miracles upon demand, but asking each of them to develop a major distributed computing system seems like one miracle too many. At a minimum, it would be unnecessarily costly.

Moreover, some portions of the bill *do* call for communication from the precincts to a statewide server. In that context, it would be utterly implausible for each county to choose its own encryption method, as Section 11 seems to call for. Surely the authors of the bill did not want to ask the secretary of state's office to adapt the centralized system to each of 87 different secure communication protocols. If the legislature rushes this bill through with this muddle still in place, the election officials at the local and state level will have to try their best to sort it out, further adding to their burden.

We were also struck by how detailed some of the specifications are. We are not convinced that members of the legislature have the relevant systems development expertise to be able to judge the risk that split VPN tunnels pose to security or the appropriateness of a 5-second time limit for updates (as opposed to a 1.5-second limit for retrieval). Many of these details may in fact be quite reasonable; but if any one of them proves not to be, a return trip through the legislative process will be necessary. It would seem more prudent to charge the secretary of state with developing the necessary technical standards.

With regard to the 5-second and 1.5-second limits (contained in Section 12), we are also concerned by two ways in which these performance standards fail to comply with the industry's best practices for how performance requirements ought to be written. These two issues are of concern both as potential sources of confusion and difficulty during system development and also as a sign that the bill may not have been adequately vetted by professionals. As such, we wonder what other unpleasant surprises would surface during implementation.

One problem with the performance standards is that they fail to clearly communicate what aspect of system performance is being constrained. The language in Section 12 starts out as follows: "The electronic roster system connection must contain enough bandwidth to handle the processing time …." The technical term "bandwidth" is a measurement relevant to a communication channel, indicating the rate at which information can flow over the channel. As such the juxtaposition of "bandwidth" with "connection" makes good sense; but then we get to "processing time," which is a totally different kind of measure: it refers to how much time elapses while the server (at the other end of the connection) processes a transaction, prior to replying. Reading onward in Section 12, to the portion where the specific 5-second and 1.5-second limits arise, we are left with the distinct impression that neither bandwidth nor processing time is in fact the intended target of the specification. Instead, the drafters of this language seem to be trying to specify the total response time of the system, which includes both communication latencies (related to, but distinct from, bandwidth) and processing time.

Another problem with the performance standards is that they specify only maximum times. Generally absolute maximums are only specified for critical "do or die" systems such as cardiac pacemakers. Instead, the usual practice would be to specify that 99% (or 99.9%) of transactions are completed within some time limit. Given a specification of

that form, a statistician can tell you how many trial transactions need to be tested (under realistic load) in order to be reasonably confident that any violation of the specification would have shown up.

Moreover, whether the maximum is specified rigidly or with some small wiggle room, it doesn't provide any indication of the rate at which the system would routinely be processing transactions. Hopefully, many transactions take well less than the stated time limit. Normally, a performance specification will provide some guidance in this regard, such as the average response time (under a stated load) or the total number of transactions that must be processed within a period of time.