TRANSCRIPT

ELECTORAL MATTERS COMMITTEE

Inquiry into electronic voting

Melbourne — 22 August 2016

Members

Ms Louise Asher — Chair

Ms Ros Spence — Deputy Chair

Ms Lizzie Blandthorn

Mr Martin Dixon

Mr Russell Northe

Ms Fiona Patten

Mr Adem Somyurek

Staff

Executive officer: Mr Mark Roberts

Research officer: Mr Nathaniel Reader

Witness

Dr Vanessa Teague.

Necessary corrections to be notified to
executive officer of committee

The CHAIR — Dr Teague, you were here when the introductions of the committee were done earlier, I think. Is that correct?

Dr TEAGUE — No, I missed it, I am sorry.

The CHAIR — Okay. This is our first day of hearings, and we have two days of this. We have selected people and they have been asked if they would agree to our request to flesh out a little bit more of their submissions. You have received, I understand, a copy of the guide to giving evidence at a public hearing?

Dr TEAGUE — Yes.

The CHAIR — Basically these hearings are protected by privilege, which means you can, like members of Parliament — whilst urging caution and truth and whatever — basically say what you would like in here, but the advice to you would be not to repeat it outside, particularly if it is going to involve somebody's reputation. So when you start can you please state your full name and your business address and indicate to the committee whether you are attending in a private capacity or whether you are representing an organisation. You will be aware that Hansard is here taking down the evidence. My suggestion to you, after you give your name and indicate to us whether you are appearing in your own capacity or on behalf of some organisation, is that perhaps you might want to address your submission and then we can ask some questions in the timeslot of 20 minutes. Thank you very much for being here, and I invite you to commence.

Dr TEAGUE — Thank you. I am Vanessa Teague. I am a senior lecturer in the department of computing and information systems at the University of Melbourne, so my business address is Parkville. I am appearing in a private capacity, not representing the university.

My background is in cryptography and computer security, and my area of research focus is on the verification of results in electronic elections, so I have a lot to say about iVote. I also have a little bit to say about the vVote if you are interested, but since iVote seems to be the discussion of the day, let us talk about that and the internet.

The iVote system was not secure. It did not adequately protect vote privacy. It did not defend voters against coercion, and, most importantly, it does not really provide any genuine opportunity for anyone to verify that it gets the right answers. I can expand upon all of that for as long as you feel like listening, but let us talk, I think most importantly, about verification, because it is really that that is the key question here. Are we actually getting the votes out of the electronic system that accurately match the intentions of eligible voters going in?

Now, there are two separate steps to iVote's verification mechanism. One is the opportunity for voters to telephone in and query a phone-based system to ask what vote has been recorded on their behalf. The second step is what you might call a backend verification process, where somehow the votes that come off the verification system are supposed to be reconciled with the votes that go into the count. Both of these systems are badly flawed. The main thing that is wrong with the telephone-based verification mechanism is simply that we did not hear what its failure rate was until very, very recently. I think I wrote in my submission that the New South Wales Electoral Commission had put up on their website this clear statement that says, '1.7 per cent of electors who voted using iVote used the verification service and none of them identified any anomalies with their vote'.

Now, the New South Wales equivalent of your committee met two weeks ago, and they asked the iVote manager, Mark Radcliffe, 'How many voters called in and tried to verify but failed?'. He said, 'Actually, we got 627 calls that failed to retrieve any vote at all'. Now, that is out of around 5000. That sounds to me like an error rate of about 10 per cent. There are two serious things there. One is that that is a 10 per cent verification failure rate. Two is that that was not made known to anybody until more than a year after the election. Now, we do not necessarily know what caused those verification failures. There could be all kinds of innocent explanations, and we are talking about failures of verification, not necessarily failures of voting. It could be all kinds of innocent things, it could be that people just forgot their receipt number, it could be that people got confused about their credentials or it could be that votes were dropped off the verification service but not dropped off the main voting service.

But also there are non— innocent explanations for that, and that is exactly the failure pattern that we would see if an attacker had deliberately manipulated the vote using the security problem that Alex Halderman and I identified during the run of the election, because if you had broken into somebody's web browsing session and deliberately manipulated their vote, you would not merrily give them back exactly the right receipt number that allowed them to login and retrieve the vote that you had manipulated; you would instead block the return of the proper receipt number or give them the wrong receipt number and organise for them to fail to retrieve any verification information at all. So all we really know is there is a 10 per cent verification failure rate. We do not really know why, and we do not know whether it really reflects a 10 per cent failure of the votes to be accurate or whether there is some other innocent explanation. So that is the first thing.

The second thing concerns the verification of the subsequent processing, because there are really two things we need to know if we want to ask about whether an electronic voting system gets the answer that the voters chose. Each voter needs to know that their vote went accurately into the system, and then there needs to be some kind of process by which some scrutineers or some independent parties get some kind of evidence that the subsequent processing of all those votes is correct.

Now, the New South Wales Electoral Commission's submission into the New South Wales inquiry said that there was an opportunity for some independent parties to reconcile the votes that came off the verification service with the votes that came off the core voting database and to check that those votes went accurately into the count.

The PwC report that observed that process said something very, very much less. It said that the third parties got the opportunity to verify that the votes on the verification service matched the votes that came off the core voting system, but it did not say anything about those people being able to check that they were the votes that went accurately into the count. This is an important distinction, and I do not think you need a PhD from a fancy university to see that there is a big difference between checking whether two lists of votes are the same and checking whether those two lists of votes are exactly the votes that go into the count at the end.

I think this is a question for Mr Brightwell. I think it is really important to understand who actually got to verify what and exactly what the process was, because I do not think that there was in fact any real opportunity for anybody to verify that the votes that went into that count were exactly the ones that had come off the verification service, and I think the distinction is important.

Now, there is a cryptic line in PwC's report that suggests that something did not go entirely right with this process, but the line is garbled and does not make any sense. What I do think is that the reason there was no opportunity for anybody to verify this crucial step is simply that the sets of votes were never meant to be the same, because the iVote protocol includes an opportunity for votes to be deleted at that point. Remember that you could cancel your vote and vote again; right? But then the question is: 'Okay, if I can cancel my vote and vote again, who fishes my vote out of the electronic ballot box? And what is the opportunity for anybody else to verify that this process of removing some votes out of the electronic ballot box has been done according to exactly the right process?'. The answer, I think, is that after the votes off the verification server were reconciled with the votes on the core voting system, Mr Brightwell had the opportunity to remove some votes — which he is supposed to do, if they have been cancelled and the person has revoted. But the thing that has, again, never been made clear is: how many votes were removed at that point, and who verified that Mr Brightwell removed exactly, and only, the votes that were supposed to be removed? I could talk about iVote security all day, but I think I will stop at that point and let you ask what you think is important.

Mr DIXON — I just want to get my head around verification. There seem to be different reasons for verifying and who verifies and why they verify. Can you just talk a bit more about verification, because it seems to be a key issue?

Dr TEAGUE — Yes, it does. First of all, it has got nothing to do with computers in particular. If you think about the paper-based process, I walk into a polling place, I fill in my own vote with a pencil, I can look at it and see that it accurately reflects my vote and then I fold it up and put it in a ballot box. Then we have a whole lot of established procedures for letting scrutineers observe the polling place process. You are allowed to show up at the beginning of the day and check that the box is empty and so on and so on. You are allowed to stand there and watch while they count the votes at the end of the day. That is all about verifying that the election outcome matches the expressed intention of the voters. So when we talk about verification in an electronic context, the aim is to replicate that kind of process of generating evidence that the election outcome is right.

In the case of an electronic system it usually makes sense to split the process into those two different steps, just like they are in paper: one step in which the individual voter gathers evidence that their particular vote has been accurately expressed and accurately included in the count; and then a separate, independent process in which scrutineers or the public or some kind of more global process allows somebody to check that all of the electronic votes have been properly dealt with and that the outcome accurately matches them.

This, by the way, is one of the reasons that it is an awful lot easier to do electronic verification in a polling place rather than over the internet, simply because it is harder to do. It is complicated to do in an electronic context, and the opportunity to do the process with some help makes a big difference to being able to do it properly and also because the most easy and obvious way of being able to verify if your electronic vote is cast in the way that you intended is just to print out a piece of paper that shows you what vote you have cast and then to fold that up and put it in a ballot box or put it in storage or something at the polling place, as a backup evidence trail.

Ms PATTEN — Do you think that is an effective evidence trail? I have read other submissions that have suggested that with electronic voting, as long as they have voted and then printed off a receipt — which probably raises some issues that Dr Naish remarked on — that could be used as incentive or disincentive to prove how you had voted?

Dr TEAGUE — It depends what you do with it. What I am envisaging is you print it off in a polling place, you look at it carefully and then you put it into a ballot box in the polling place. So the assumption would be people cannot walk out with it because then it would not get counted.

Ms PATTEN — So it would be automatically printed off to ensure — —

Dr TEAGUE — Yes.

The CHAIR — Can I ask you about scrutineers, because I think most members of the general public would not even know what scrutineers are, but obviously you have an appreciation of how important they are to political parties and candidates feeling secure about counts.

Dr TEAGUE — Yes.

The CHAIR — Would an electronic voting system mean that the scrutineers, who are just people who have a view about supporting a candidate or a party or whatever, would have to have some sort of computer qualifications, or would the ordinary person test still be able to apply to scrutineers if the system went totally electronic?

Dr TEAGUE — I think it depends on the system. If you were using what you might describe as the low-tech system, where the computer just prints out a paper ballot and people look at it and put it in a box, then there is not necessarily any special training or computer knowledge required to watch what happens with that paper evidence. If you are building paper evidence using fancy cryptography as in the vVote project, which is a project that I had a lot to do with and a lot to do with the design of, then there is a high degree of knowledge of computer processes and cryptography required to really understand what the evidence trail is showing.

Ms SPENCE — Obviously you have had a look at many, many electronic voting systems. Would you be able to talk to us a bit about not necessarily the iVote system, which we are familiar with, but what you see as positives or extreme negatives in other systems elsewhere?

Dr TEAGUE — I think there is a huge difference between systems that are designed for voting at a polling place and systems designed for voting over the internet. I think there are a lot of sensible solutions in the polling place, as we have discussed — printing a plain paper record, running an end-to-end verifiable system like vVote. These kinds of things do build a genuine evidence trail. I do not know of any system that builds a genuine evidence trail running over the internet, because the complexity of what you have to get people to do in order to really get evidence that their vote is going properly into the count is too hard to expect ordinary people to do.

Furthermore, you add a whole lot of other issues like trying to authenticate the voter at the other end, for example, making sure that the person at the other end of the transmission is the eligible voter that you think they are and not a bot that is submitting 50 000 votes from ripped-off credentials. So I do not think there are any good solutions for running over the internet. I think there are a variety of sensible solutions in a polling place, as long as you have worked to build the evidence trail that supports the election result.

Ms BLANDTHORN — My question was similar. With your extensive experience, what type of system, in your mind, would be the best type of move towards electronic voting, either in part or full for the process from go to whoa?

Dr TEAGUE — I did a lot of work on the vVote system. I recognise that the VEC has decided that that was too complicated, and I respect that. So I would say that, except that there were obviously some difficulties in actually running voters through the verification steps that they really needed to run. If that was too hard, then my suggestion would be simply using all that open source, using the interface design that they have already got and just printing out a plain-paper record at a polling place and then just either counting that manually or using it as a backup evidence trail for auditing.

The CHAIR — Could I ask you about overseas and interstate voters because, again from a practical point of view, they are the voters who are most impacted on by — I am probably offending — the decline in services from Australia Post and this convoluted system to try and get a vote. I am very focused on: how we look at this mass of voters that have to go through a ridiculous process at the moment to get their right to vote if they happen to be holidaying. Those numbers are increasing and increasing, which is why the London booth was targeted for electronic voting.

Dr TEAGUE — Right.

The CHAIR — In Australian terms the second-biggest booth is Hong Kong. I am told that the AEC actually sends staff to the embassy to process votes. Again, when you look at the cost of that, it is a horrendous process to try and give people their right to vote in an election. So my inclination is to think, ‘Can we find a system that could deal with this mass of overseas votes?’. I do not want to put words into your mouth, but are you telling me that an internet voting system is fraught with difficulty for those votes and it is going to be inherently unreliable?

Dr TEAGUE — Yes, indeed. I think you can deliver candidate information online. I do not see any reason for posting out a piece of paper. That seems like you are posting out public information that they could look up online. But I do not see that returning an electronic vote over the internet is feasible in a way that gives you evidence that it was returned in the way that the voter intended.

The CHAIR — Could you send a ballot paper over the internet?

Dr TEAGUE — You could send blank information, is my point. But in terms of returning how the person voted without the opportunity for that to be manipulated online I think — sending it over the internet is not going to be as reliable as post, even given the serious problems with the reliability of post. I think that is the thing — the argument that the postal service is declining is a serious issue. It is a problem. It is not entirely reliable, but that is not a reason to cross across to a system that is even less reliable and, if anything, even more likely to fail in a way that is silent. At least the opportunity for one person to manipulate a very large number of postal votes without detection is pretty limited, whereas what we showed in New South Wales is that the opportunity for one person to manipulate a very large number of electronic votes is entirely practical.

Ms BLANDTHORN — Does the same goes for, say, the internet voting that was available to some elements of the defence forces and whatnot? It might be a technically naive question, but does the same apply?

Dr TEAGUE — The same thing. Yes.

The CHAIR — Even with defence security, which I think one of the submissions has addressed, that they had access to defence department computer technology in the processing of those votes. So that submission has given me the impression that the security on those votes is in advance of the security on other votes. But even with that, you still think it is able to be manipulated?

Dr TEAGUE — I read what was available of the security analysis of the federal ADF Internet voting trial of 2007. There was some kind of consultant paid to write some kind of a report — I would have to look up exactly who it was about the security of that system, and the extent of the security analysis was entirely about whether running this thing on that network impacted the security of defence's network. There was very little about actually analysing the security or verifiability of those votes.

Remember that when we are talking about elections, we are talking about something in which anybody has an opinion. We are not necessarily talking about the Chinese breaking in from the other side of the world, although that is an issue as well. We are also talking about the capacity for the system administrator on a network to influence other people's votes. I would not be comfortable with a voting system in which one person had the opportunity to influence the votes of the other people in their barracks or certainly not the opportunity to change them without detection. That system, which also had no meaningful verification, certainly did not defend against the opportunity for somebody at the wrong point with control over somebody else's communications to change their votes.

I know what I need to say about this; I do not know how much of the geek news you read, but just on the subject of whether the defence department's network is perfectly secure, I do not know if you saw — this is probably not the sort of geek news that you read — that somebody dumped very recently out onto the public internet a vast collection of NSA intrusion software from about 2013. Initially it was not clear whether it really was NSA stuff, but the word seems to be at the moment that even the US NSA cannot keep their special methods of breaking into other people's computers secret. It was not just explanations of the methods; it was code for doing the stuff. So I do not believe the argument that the defence network is perfectly secure from external attack, because if the NSA cannot keep their ways of breaking into other people's computers secret, then I do not believe the ADF can either.

Ms PATTEN — I just have a question. I have been looking at the vVote system, which you and other experts are suggesting that with kind of off-line electronic voting there is some common sense about streamlining some of that voting system — and I think also given the length of pre-polling these days and the number of people who are accessing polling booths over a period of time. Why do you think only 1300 people took advantage of vVote? Was that the limit of the cohort that we allowed, or would you agree with the VEC that it was too complicated?

Dr TEAGUE — Probably both. I did not have anything to do with actually rolling it out, so I do not really know. Ask Craig Burton, I think. But I think the answer is probably both. Probably it is complicated and hard for people to understand. Probably it was complicated for the electoral commission to understand, which is something I did not really appreciate through the early phases of the process, and I think they were nervous about it because it was complicated for them. So I think they probably did not advertise it or market it in a way that they would have if it had been simpler and they had had a better understanding and more confidence that they could actually run it. They did not advertise all of the verification steps that voters should have been told to do, and I think again part of that is just, first of all, not really understanding what they were for and, secondly, just not wanting to make it any harder than it seemed to have to be. I do take on board that that was too complicated. I still think that as a basic approach to the future of voting that kind of approach is really the only sensible way to go — that voting has to be transparent; it has to have openly available source code; and it has to be building an evidence trail. I just do not think there is any future in running a proprietary system where you cannot see what is going on and you do not get any evidence at the end of the election that the votes that come out of it were the votes that went into it.

Ms PATTEN — That was obviously an off-line system, so what about the notion of enabling people to check off on the electoral roll at the polling booth in an electronic form that stops someone from voting early and voting often at other polling booths?

Dr TEAGUE — So two things; just let me take one step back. It actually was not an off-line system. It was on a virtual private network, but it was in a polling place. And then as to your question of what about automatic roll mark-off, I think the VEC already does that, but I think you would have to ask them. I do not know exactly what they do. But I think the short answer about that is back to this idea of whether if something goes wrong you are going to find out about it, which is really in some ways the key point here. If something goes wrong with this system, are you going to know? If nothing else, that is the question to ask I think. Am I entrusting the integrity of my election to something which, if there is a malfunction or a security problem, I will find out, or am I entrusting it to a process in which there is a possibility for undetected problems to change the outcome without anybody knowing? It seems to me that electronic roll mark-off fits fine into the category of things which if something goes wrong, you are going to find out. It is not to say that things cannot go wrong; it is that you are not going to get widespread undetected fraud because somebody broke into the system and did something naughty.

Mr DIXON — When you look at information technology, where it has come in the last 10 years, and say, ‘Well, what is it going to be like in 10 years?’, do you think we are going to get to the stage where all of the bad guys are going to be as advanced in 10 years and it is going to be a zero-sum game still?

Dr TEAGUE — I think the sophistication of attacks we see on the internet is growing probably faster than our ability to defend against them. As you say, it is almost an arms race. We actually see less DDoS than we did 10 years ago, but then we see more targeted, specific intrusions into particular pieces of infrastructure. We see different things. I do think there are sensible solutions in the polling place, mostly designed around building an evidence trail that does not rely on trusting the software system to be secure. I think that is the thing. Building an evidence trail that does not rely on trusting the computer system to be secure is really the key to making it okay.

The CHAIR — Thank you so much, Dr Teague, for coming along. It is certainly a fascinating inquiry. We greatly appreciate the detail in your submission and your comments, and in fact we have used the extra 5 minutes, so thank you for your willingness to extend that participation in our inquiry. You will receive a copy of the Hansard transcript in about two weeks or so, and again you are free to correct any errors in the transcript but obviously not to change the content of what you have said. Thank you very, very much for participating in the committee’s inquiry. It is greatly appreciated.

Dr TEAGUE — Thank you. I am easy to find; I am just up the hill if you want to ask any more questions.

Witness withdrew.