Submission to the EMC Inquiry in to E---Voting (1 July 2016)

I would like to submit the following I have learned from designing, developing and deploying electronic voting (including Internet based voting systems) over almost twenty years. I have written in a non---technical style and have attempted to make several themes in election security easy to read and understand.

I am an Australian, my name is Craig Burton and I provided e---voting systems in the US, the UK and in federal and state elections in Australia from 1998. I am trained in Computer Science and have worked in computer security for many years. Last year I finished a six---year contract with the Victorian Electoral Commission (VEC) to build and deploy voting systems, the result of which was the first use world wide of the Prêt à Voter e---voting system.

I now only deploy *verifiable poll place voting systems*. I make a personal appeal below. I then expand on this summary at length for your interest and information:

- Automation is good, but not everything automates well. Some processes cannot be automated at all, such as driving tests. Some important government processes will always require citizens to attend a controlled location and be manually supervised. Voting is likely to remain one of them.
- It is hard for a single person to affect a paper election outcome. On the other hand it is too easy for a single developer or operator to interfere in (or halt) an automated election. Computer issues during elections should not be treated like simple admin issues because a small technical hitch can have dire implications.
- Physical world risks and computer risks cannot be compared, even if it looks like they can. Errors in manual elections are usually random and do not favour candidates. Errors or fraud in software can systemically (not randomly) damage elections.
- E---banking bears little resemblance to Internet voting in its security design
- The electoral commission should never be in a position where it is challenged and cannot issue convincing proof of an election outcome. Where computers are involved, this requires new permanent, in---house technical security capabilities. Security modernisation is already needed for current computer systems that handle votes such as keying---in facilities and vote---counting software.
- Elections cannot be outsourced, but Internet voting and highly technical automation effectively approach this. Automated elections must remain human intensive projects and not "set---and---forget" exercises.
- Like the UK, Australia needs the role of "independent observer" in its electoral Acts to allow more oversight of technical systems, by "anyone".
- A promising way forward for computers in voting has been demonstrated successfully in Victoria with the verifiable system vVote. This system is now well known in e---voting research around the world. This needs further investment.

My personal plea to EMC

It is surprising to many people that Internet voting has been used in Australian state and federal elections already. Also, not many Australians know about accessible e---voting systems or of the work of the VEC. This is largely due to these electronic voting systems being deployed for certain classes of elector such as military, those with low vision or blindness, inability to read in English or any language and those who cannot hold or move a pencil due to motor problems or injury. This has been my work and I have provided all of these systems. However, my experiences have led me to change my mind about the introduction of computers in to vote casting.

I still see many opportunities for the use of computers in election processes. Computers are already heavily relied on in commissions for printing, finance, administration, logistics, promotions and customer relationship management. Computers can be further deployed effectively and safely. They have the potential to increase the accuracy and speed of the count and improve accessibility. I know this because I have deployed the VEC's system vVote that does these things and can even provide a high integrity election outcome if it has been compromised ("hacked") or has experienced a serious IT failure. I explain this quite outlandish claim below.

However, I want to first say that some uses of computers in elections should not be pursued – the most dangerous of all is Internet voting. There is so much naïve support for this I feel I have to assert why it should in fact be discontinued world wide wherever high stakes public elections use it (which is not many places!).

I founded an Internet voting company with KPMG Consulting Australia in 2001. It was spun out in 2003 and ran many Internet based elections for basic polls right up to the 2007 Australian federal election and the 2008 DNC selection of Barack Obama. The same system now runs the Academy Awards ("The Oscars") selection. The NSW iVote system's first run in 2011 was on my company's software. The first cryptographic ballot system run by the current NSW iVote supplier, Scytl SA was run on top of my system, in 2002. This was a promising area of Internet automation after the Tech Boom of 2000. Some very bright people have worked on this (including some former NASA staff who worked at VoteHere.net) and hundreds of millions of dollars of investment and procurement has taken place with Internet voting systems. High integrity, safe, Internet voting however, turns out to be an intractable problem.

Since 2011 I have been directly *opposing* the use of Internet voting in any high stakes elections. This is because I have now occupied roles of system provider, electoral commission staff, security analyst and voter and have a well rounded picture of the *total risk surface* of elections. I have seen the growing panoply of Internet borne crime and I have read books with titles like *Insider Threats* which catalogue that about half of computer crimes come from within organisations. I have directly experienced that software *bugs* are a serious risk to elections. I agree with the writings of many academic specialists on computer risks to

elections. These experts include Vanessa Teague, Steven Wagner, Arvi Rubin, Alex Halderman, Richard Buckland and many more.

Still, I am confident 99% of the world's elections can actually be run over the Internet --- just *NOT the high stakes public elections*. The example above of the Oscars is an election that *can* be run via the Internet because while it would be reprehensible if there were problems with this system, an unexpected nominee winning Best Actor would not, as such, undermine public confidence in the state or national governments.

It isn't hard to see that with the Oscars chosen over the Internet, and half of Internet users accessing e---banking that there is understandable confusion among the general populace, and there is the expectation that they should be able to vote from home over the Internet. It seems that voting over the Internet is "inevitable" and that somehow we are backward for not having provided this. We are most certainly not backward – Victoria was the birthplace of the main method used world wide for paper voting. Our Electoral Commission is progressive, makes advanced use of computers, is highly capable and is very well staffed and skilled. It has resisted Internet voting for good reason.

The vast majority of information technology specialists in the world who follow or contribute to elections or election technology *do not* support Internet voting as a means of collecting votes in major elections. It is true you will find some support for this proposed method among IT people in general and even among IT people who work in information security. However, it takes an understanding of elections and preferably experience in executing elections to really see what the experts are concerned about. When I explain election legal requirements and election risks to IT professionals they become nervous about Internet voting.

My advice is that Victoria should pursue a number of new technologies in its State and Council elections to make the process run efficiently and quickly. Victoria (and everywhere else) should not use remote, unsupervised Internet voting systems, or should cease using it. The decision as to what technologies to use and what to avoid should be taken very carefully, with consultation and with transparency. There is no rush to do this.

That is the end of my personal appeal to EMC. In the rest of this submission I will expand at length on the various risks present in high stakes elections, along with new risks introduced by the use of computers for your information or interest.

My work since 2012 has been in what is known as *verifiable electronic voting* systems. Since 2015 I am working in a part time capacity running coasca.com, a British start---up that maintains a public open source version of VEC's vVote (also called coasca). This version of the software can be downloaded and used by anyone for free and indeed has been downloaded and is being tested at several sites around the world.

Not everything automates well

There is no way the world is going back to yellow envelope inter office mail, nor The Trading Post newspaper classifieds. These have automated well. However, there is a minority of activities that computers have not automated nearly as well, an example being credit card authorisation. Credit card Internet payments are common but this system has been attacked successfully for many years and even with chip and PIN, credit cards continue to be the most simple and lucrative source of computer crime. If credit cards were invented tomorrow as a way of doing Internet payments they would probably be banned. This is somewhat like motorbikes because it has been said that if motorbikes were invented tomorrow they would never be allowed on the road due to present regulations over airbags and myriad car safety requirements. Credit cards and motorbikes are examples of technology grandfathering. Really, for a safer world they just would not exist. However, we have these technologies now and we are pushing on with them, even though they are risky. Both banks and road traffic authorities have to take out expensive insurance for these risks.

I propose a second class of automation --- processes that have not automated well at all and have been discontinued --- and I'm putting Internet voting (that is, voting from home on your PC or ipad) in this category. Why can't we just take e---banking security or web---based polling systems and use them here in Australia for Internet voting? Why isn't Internet voting in my first category of "established" technology that "is not perfect" but that we must tolerate? The answer is that Internet voting has not worked well or not worked at all when one considers the *total risks involved*. And, unlike credit cards or motor cycles which must be respectively integrated in to banking and road safety ongoing because they are already out there, there is no imperative for the world to automate its voting systems to run over the Internet.

Internet voting was first used in 1995, around the same time e---banking was first deployed. E---banking is everywhere. Internet voting is almost nowhere in major elections. Trials of Internet voting in the UK, France, Norway, the US and other places have been halted. Internet voting introduced in Hawaii caused an 80% drop in attendance. Internet voting in NSW publicly suffered bugs (2011) and vulnerabilities (2015). I predict things are going to get worse, not better for Internet voting. But I think things will get better for e---banking. I explain below why e---banking will never provide any kind of basis for Internet voting.

Internet voting is not alone in my category of non---starters. E---conveyancing has cost homeowners their houses. Nigerian scammers sold two houses over the web (located in Perth) out from underneath their owners, with the unwitting cooperation of the real estate agent. E---conveyancing does not work. Voting on computers from home does not work. In all cases of these e---technologies, the proponents need turn up in person to sell a house, or to vote. Some important processes require attendance, trained supervision and physical security.

My last category is human processes that will never automate. Driver license practical testing will not automate, nor has it even been attempted. This is a

good thing even though video games would seem to provide a very immersive home driving test with all sorts of driving conditions ready to challenge the student driver. I predict that we will never be able to sit a driving test over the Internet and print out a driver license at home. I know this because it is telling that student pilots have to use a flight simulator for hundreds of hours but not one minute in the simulator counts for flying time toward actually getting the pilot license. This is again because the student needs to attend to drive a real car or fly a real plane.

In the second category of grudgingly accepted technology I put polling place e---voting. Voting computers that replaced lever voting machines (grandfathering again) are well known. These machines record or transmit votes electronically – Direct Recording and Enumeration (DRE). Two thirds of US states have mothballed their DRE machines and now vote on paper again. All the DRE machines in Ireland were scrapped at a cost of E60m.

Those thousands of poll place voting machines which have been discarded were probably not complicit in large scale fraud. One important issue was reliability. In elections, we need systems to work. This is why we use pencils not pens – pencils work when they are wet and they can write on wet paper. Automating the completion and casting of a ballot could have gone pretty well, but it took only a few failures to highlight that these machines can mis---record votes and that this may be undetectable. In elections, when something goes wrong, we always need to know.

Making computers able to always demonstrate when something is wrong has become the focus of election systems research. This is a new approach from trying to make e---voting systems "flawless" or attempting to add more and more security to defend voting systems. Neither of these approaches can always demonstrate when something fails and, voting computers just end up in the same security arms race as desktop computers. Instead, we do not place any implicit trust in computers used for voting. We assume from the start they are going to have flaws, can't be trusted, and may well fail in some way. We then set up a new verifiability system to catch them out. I call this verifiability approach a prop. A prop is a low---tech and simple solution to holding something up, as opposed to some new incomprehensible black magic technology. We are setting up voting systems to fail, on purpose.

I now explain why effectively "setting up voting systems for failure" is not as bad as it sounds and it lays a much better foundation for detecting and managing risks.

Elections are risky – computers magnify these risks

The paper election process admits a number of incumbent risks. These risks affect ballot papers and they can affect the reported public outcome of the election. When the election race is very close, a losing candidate may request a re---count to confirm the race result. The historical differences in results between counts and recounts (or internal re---checks) in Australian elections is about

0.05% (and for certain races this can be higher). This seems small but in fact many races are declared on a narrower margin than this. Why would an electoral authority be comfortable with such risks and such an error rate?

Imagine an election race has 1000 votes for two candidates. We count them. Jane has 505 votes and Bill has 495. It's going to be a pretty close win for Jane (10 vote margin). Paper voting is not 100% accurate, because each marked paper ballot is handled differently, by different hands at different times – introducing an imprecision engineers would call "noise". In the absence of a report of a sizeable integrity breach, batches of votes in doubt or missing are most likely due to random errors or random damage, and not votes mishandled or discarded on purpose because, specifically, they are Jane votes.

The total error is the total effect of small mistakes and events. Nonetheless if some votes are in doubt you might think that there's a good chance Bill might now win. In fact, if we randomly remove 20 votes Jane still has a 99.4% chance of winning. How did I work out that 20 missing votes still sees Jane so certain to win in such a close race? The chance of Bill winning due to omissions is that somehow 16 or more of the 20 missing votes is a Jane vote. This tips the balance back to Jane less than 490 and Bill more than 490. What's the chance of 20 votes removed at random being 16 or more Jane votes - it is about 1 in 170 (for probability wonks: ((20,16)+(20,17)+(20,18)+(20,19)+(20,20))/2^20).

So you can see the system is highly robust to small errors, loss, damage and even fraud affecting small random numbers of votes, even when the total number of impacted votes from lots of little problems could be quite a large number. Random errors do not favour certain candidates.

The Australian electoral process is carefully designed so that loss, damage and crime are detected by a number of cross---checks on batches of 50 or fewer votes. The larger an effect on the integrity of votes, the more likely a cross---check picks this up. When 1375 votes went missing in West Australia this was in fact a *triumph for robust design* for at least two reasons: firstly the loss was actually detected at all, and secondly it was accurately determined the missing ballots *would affect the winning margin* of one or more candidates. And so the race was run again. In an inferior design loss or fraud occur undetected and the wrong people win, then the missing votes might turn up at a later date!

When electronic voting is introduced, the way votes are handled does not resemble the paper voting process at all. Electronic systems do not handle individual votes differently in the sense that in place of many hands touching the votes, one machine algorithm touches all of the votes. Immediately it can be seen that a flaw in the voting algorithm affects all the votes whereas one clumsy election worker may affect only a proportion of the votes.

The greatest risk for electronic voting is bugs. Obviously good software engineering is important to reduce the number of bugs but in the end, bugs can occur anywhere, and it is never possible to know all the bugs in a piece of software have been found. Why is this the case? Because the very same random indeed human, mishandling of votes also results in mistakes in software.

It isn't possible to get all the bugs out of software. Here is an analogy: it's easy to find something wrong in a mechanical clock because it has only 11 interacting parts and it can be observed doing the same thing all day long. Mechanical clocks can be made hyper accurate so they slip one second in twenty years. In contrast, a computer has 100 million interacting parts (lines of software) and may never do exactly the same sequence of things twice. Considerable cross---checks, audits, testing and human vigilance are needed to make computers reliable. Aeroplanes have three entire computers for every task in case two fail, and military devices have multiple computers programmed by different companies, set up to all work together attempting to agree on each and every step in a high risk computation. If any systems produce different results from the same data, alarms go off.

So introducing computers to handle votes introduces a new kind of risk: the risk that a bug will affect the votes it handles and that no one will be able to see this. It is not likely a computer bug will affect only 20 votes in 1000. What's worse, in the paper voting world, a misconfigured letter opener which chops the bottom margin off postal ballot papers may still not damage the electors' voting marks. In a computer the data are the votes, so damage to the data is likely damage to the votes. Since Jane vote data are different to Bill vote data you can see that a bug may well cause more problems for one kind of vote than another by affecting the same bits of each vote, perhaps the Jane bits. This is called a systemic error because it will not randomly affect Jane and Bill votes with equal likelihood.

It goes without saying that cyber crime that seeks to do more than simply disrupt elections is most likely to introduce systemic influence on the election outcome. Again, since one algorithm touches all the votes, an attacker does not need many hands – two hands can touch all the votes. You might expect an entire database of Jane votes changing to Bill votes would be a pretty obvious sign of fraud. Above you read that randomly changing votes has almost no effect on the outcome. However, systemically changing 16 votes in 1000 would change the outcome in our example.

Computers magnify these risks. It is better to assume the risks are present than to design for perfect systems or attempt to try to control all the risks. In Victoria in 2014, Prahran was won by a margin of 31 votes and vVote took 42 votes there. Being able to prove all's well becomes very important, even with small uses of a voting system. It should be said that if Internet voting or an unverifiable system is used to collect e---votes, then it becomes the weak link and all the physical security carefully executed in poll places may be for naught: the election is compromised via the e---voting.

Why can't an e---banking system do Internet voting?

This is a good question because e---banking has been made very strong and reliable given it handles trillions of dollars. Doesn't this mean that Internet voting should be a solved problem? To illustrate why not, I provide another analogy.

If the AEC was a bank that provided Internet banking, *but with the rules of elections*, you would see this: The AEC could take a deposit from you, add it to the bank's vault but it could *not* record how much money you deposited *against your name*. Or, the AEC could take a deposit of ten dollars, it could record that *you* made a deposit, but it could not record the ten dollar deposit *was from you*. You would get no end---of---month bank statement of your transactions. You would not be able to "undo" an errant transaction. You would not be able to know your balance had changed because someone had transferred money from your account without your permission. The AEC can't keep a running total of all the money it has. This would not be a very effective banking system – and yet, this is exactly how elections capture and record anonymous votes. I hope you can see that all the brilliant work on financial transactions, which record who, how much, when, and where from, are broken by our requirements to collect and store votes anonymously. This is but one reason why e---banking does not mean that Internet voting is a solved problem.

Paper balloting has had 150 years to perfect this recording process – 150 years of "debugging". It is cleverly designed so that no single actor or no few---actors are likely to get away with attempting to affect the electoral outcome. Instead, it would take very many actors, working in collusion, to bring about a large change in the result and have any chance of getting away with it.

In contrast, a single actor with a single computer can enter and change data in a "secure" system. There are reports of such accesses, thefts and data corruption every week. It is easier for criminals to do this if the system provides no proof back to its users that their data is safe. And exactly this has been the design of e---voting for a long time – voting computers turn votes in to electrons and use standard IT techniques to protect them. Voting computers offer no proof that a vote made it all the way to the count unchanged. If the computer could offer this proof to the elector, it would have to do it in a way that the elector could not show someone else – so this is a hard proof to design.

Prove it

For the first time, it is now possible to collect votes electronically in high---stakes elections and know, with great certainty, that all has gone well – or as the case may be, that there has been trouble. In all circumstances, as with West Australia, it is crucial to know the integrity of the election result. End to end verifiable electronic voting (E2EV) as it is called, has introduced this.

There are two systems in the world that can do it. One of them – vVote --- has been used in a State election in Australia (2014, the other is called ScanTegrity and has been used in the US). vVote is open source software, it is the subject of a number of peer---reviewed journal articles and it is held in high regard by the same people who are deeply critical of the automation of elections. The key difference those experts defer to in their confidence is that vVote is a polling place E2EV system. It is connected to the Internet, but the voting devices are controlled by the VEC and they are operated by staff specially trained to support electors using those devices. How can these computers connected to the Internet be more trustworthy than your iPad at home? This is because a disposable printed receipt system that operates with the computers provides surety that the system is behaving.

There are other receipting e---voting systems in the world and rather a lot of hubris and trust is heaped on the humble printed receipt. The difference in Victoria is that *the receipt is the vote*. It is a printed list of the elector's actual preferences that the elector can check and take home with them. At home they can check the list online with the VEC record to see it is present and it is unchanged. They can do this with help from others or with mobile devices without risk that anyone else can learn their vote. This is because the receipt lists their preferences in a shuffled order, different to the legal ballot order, and an order different from other voters who attended the voting location. The voter themselves had a brief opportunity in the polling place when they can see or hear that their preferences, even though they were shuffled, still reflect their choices. Since the preferences are shuffled and the ballot candidate names are removed, no actual information about the vote can be ascertained. In the end, the VEC has the ability to "un---shuffle" each receipt back to the legal ballot order so that it can count the preferences against the correct candidates.

This design is so simple, and yet so strong that a fully compromised or malfunctioning system can still deliver a correct election outcome so long as a proportion of electors confirm the receipts are right. Of course there are other things to check and vVote supports other challenges all the way from ballot creation to the final emission of the countable votes.

For clarity and for a technical reader, I would like to reinforce that verifiable voting is a security system that can always emit proof. This is quite different to a security device like a firewall. A firewall blocks attempts of unauthorised people getting in to a part of a private computer network. The firewall reports all the attempted intrusions. What it can't do is report the unauthorised people who got in somehow anyway. These successful intrusions may not be recorded anywhere because the intruders found a vulnerability in the firewall. Verifiable systems assume from the outset that intruders *will always get in if they want to*, and how do we demonstrate integrity (that our data assets are undamaged) anyway?

vVote still needs to be run on reliable computers and it still needs good staff to support it, but it is a system unlike most others – not hack proof, but able to always emit meaningful proof of integrity. If it can't emit convincing proof

something's wrong. And in elections, when something's wrong everyone needs to know.

When there is trouble in e---voting it's not good enough to issue gobbledegook technical proof that all's well because only experts can understand it. vVote provides at least some of its proofs in forms the general public can examine – and they have. When vVote ran in 2014, more than half of users checked their receipts – even people who were blind or who could not read in English. Anyone was able to visit the VEC website and review all the receipts online, check system signatures on the receipts and perform a number of other checks to their satisfaction. If they did have technical skills, they could go further and indeed a small number did download all the emissions of the system and its source code.

It only takes one person to complain that a receipt was wrong, or that a single check failed to identify an integrity breach, and they did not --- so there was very likely no breach. The general public are the auditors! This is quite different from experts poring over system records trying to find out why some anomaly occurred and whether it may have impacted vote data or not. And there are always anomalies – see the recent report from The New South Wales Electoral Commission on the security of the 2015 iVote run. Did reported anomalies affect the election? Can we independently know all's well?

So why not put vVote on the Internet?

In fact, verifiable voting makes unsupervised voting from home *worse*. Voting systems should not leave an elector with proof of their vote. At least postal voting and Internet voting provide no way to meaningfully check that your vote made it, they also don't leave you with demonstrable proof. Verifiable voting does – used from home it would issue binding proof of your vote. This could then be used to appease a coercer or to sell votes. There have already been illegal schemes in other countries to perform vote---swapping (or vote pairing) which would be made much more effective by strong proof of the vote, to name but a few new economies that would grow.

Secondly, guaranteeing privacy and secrecy when the elector is in their own home is a very hard problem and has been unsolved since postal voting was invented. Postal voting has gone on to attract a range of crimes and is often the voting catchment most maligned in a very close election. Automating postal voting would not remove these risks.

Who are you going to call to run the election?

The more advanced that election technology gets, the fewer specialists understand it, and so the integrity of the election rests on very few heads. This seems unavoidable if one is to use high technology – but it is not acceptable in elections unless there is also a way to independently measure integrity. The authority needs to measure it, the public needs to measure it and the losing candidates need to measure it. It does not suffice to ask the technical experts to tell you if they did everything absolutely right.

In the end the electoral authority cannot outsource the electoral process, by law, and if there is any trouble it is the public trust of the electoral authority that is ultimately at stake. This trust will be maintained or restored if the electoral authority can issue convincing proof that in a crisis, in fact all was well. It will not be able to do this when it relies on a third party to print reports because the electoral authority will have to defer to the expert assessment of the supplier.

For any technical or security challenge that comes up in IT there seems to be a technical solution. When a critic or academic points out a vulnerability or some design flaw, the current solution is to add a new function or put in a firewall and so on. However, in the end, the IT system is very large and complicated. Fewer and fewer people understood the whole system until finally only one person understands it all, then not even them. I hope you can see that this direction is thus wrong – a larger, more brittle, opaque and complex machine understood by fewer and fewer people. In fact we need the opposite – a simple system which can be understood by more and more people.

As an outsourced election provider, I believe that executing elections has put me and my family at risk. Being one of very few people who understood and had access to the entire Internet voting system placed me in a uniquely powerful but extremely risky position. Any voting system reliant on deep trust of one--- or few--- people or machines is vulnerable in this way.

Instead electoral authorities need native, in---house skills for computer security and e---voting. It is already common for security audits to take place, and it is the case in---house election software developers are trained in security but this is not adequate. Authorities already have highly skilled people on staff called psephologists for the minutiae of paper election challenges such as the number of decimal points to round off in a carry, how to manage a three---way tie and so on. If electoral authorities want to use computers to collect, hold, and count votes they need the capacity to understand and control integrity of such systems on staff. Security is not an audit---to---audit challenge (like accountancy!) it is a human process needed from design to retirement of systems, and it is needed 24/7 during critical deployments.

Security and reliability failures during elections need to be treated differently than other IT issues. The organic growth of IT in most institutions means they are critically IT dependent as most electoral commissions are. When computers are used in the actual live election, however, IT problems are generally handled the same way as when there is some failure in the organisation administrative system. This is not appropriate for live elections because IT failures affect access to the vote and can impact election integrity. It should be the case that IT failures are handled more like floods and power outages, “acts of god”. These events cause entire poll places to be moved and they rightly fire off small armies of staff and volunteers to right the situation and keep the election going. When there is an IT failure, it usually falls to one technician to do their best and resolve it any way possible, quickly. Instead, any IT failure affecting voters or votes has to be handled by several hands at once, with scrutiny and with full reporting and

transparency. If it is a security issue, there are mission critical ways to respond established in other fields such as military and aerospace. Elections are this important, or put another way, computers are so powerful as to be a large menace.

Special security skills are needed in---house at electoral commissions already. Electoral commissions e---count votes. This exposes votes to many of the same risks present in e---voting. To catch out errors in keying, there are spot audits of digital votes. Unfortunately, these will not catch out database errors, counting software errors or errors in backup, restore or transmissions of vote data. A modern approach to election security will shore up these systems. As it happens, verifiable voting solves these storage problems as a side effect of protecting vote casting because it can emit digital votes that anyone can download and count if they want more surety of the official result. Again, the public are the auditors.

The public are the auditors and who are the observers?

In the UK the Electoral Commission can receive applications by anyone to become an Independent Observer for an election. There is some vetting of known troublemakers but in effect any citizen can become an observer without being a representative of a candidate, as is the case in Australia. As a supplier in the UK I have hosted observers to access election computers I set up either at the electoral authority or at a data centre. I was obliged to answer any and all technical questions of technical observers and there were many.

VEC has been able to offer this role by proxy – that of a technical election volunteer and such volunteers were vital to my work at the VEC. However, it would be much better if the Electoral Act recognised these people who, after all, want to commit their time and expertise for free to make Victorian elections better. Technical observers could be present at key transactions in running an e---voting system and the observer could or should be free to report to Parliament or the Court of Disputed Returns.

The future

VEC's deployment of vVote was successful, but it is a new system and it should be improved. Making it open source means it will be used and improved by others, rather than being dragged out every four years and revivified. However, for Victorian elections, vVote is a finished work ready to be re---used and extended. While it was only used for a small catchment of electors, I believe its use can be expanded but that VEC needs support to both use it, understand it and to have it fully supported by its poll staff. vVote is a big leap ahead for elections – it will naturally take more work to fully realise its benefits, to balance e---voting and existing paper voting workloads and to make verifiable voting a native function of VEC and hopefully other electoral commissions.

There are other opportunities both within and outside vVote. vVote can provide a ballot to vote on where candidate names are randomised. This would remove donkey votes and layout biases. It would remove the need for a ballot draw.

Early voting is very popular and should probably be expanded. This would allow vVote to be served to many more people over the two weeks before election day when it is more convenient.

Even with vVote, the challenge of providing high integrity e---voting is a large one. vVote is a complex system. It is harder to set up and harder to run than Internet voting. Further, it measures and reports all sorts of influences on the data it collects – even harmless or minor errors (the "dirty washing") of manual elections are outed by the system.

Polling place staff supporting vVote do not have to be computer experts. This is important, because poll place staff are the front line "sales staff" who will, if confident and competent, guide electors to the vVote devices in preference to giving them an assisted paper vote. Those supporting vVote, however, have to provide a new kind of service to electors that is unlike others in elections.

Paper elections run like a kind of military exercise. Supplies and services are put in place in the right quantities that will be relied on and consumed under controlled conditions. vVote relies on voter audits. These are optional. However, as described above, the ballot check appears as part of e---voting and is quick and easy and we rely on a good number of electors attempting this. We also rely on an another voter---initiated audit when the vVote printed candidate list is created. In 2014 this was largely not offered to electors. The reason for this was that there were concerns that electors would stall queues getting these audits, or that commission staff would be unable to explain the audit, or handle a possible failure (an audit that detects a problem). We determined that the absence of this audit weakened vVote, but not critically because all other kinds of audits were present and were performed by many people.

Providing an election service which may or may not be used but has to be trained for and offered is quite different and new. Having to offer a system which may attract potentially difficult technical questions from electors is threatening to poll staff. vVote itself had many plugs and wires and suffered from some unrelated IT and network issues which made some poll place staff nervous about offering it to electors. vVote had to be supported in addition to all the requirements of paper voting and electronic voter lookup.

These new challenges can be overcome. There are analogies for all of these problems in conventional elections already: questions about STV, problems with seals, deliveries, eligibility, comprehension and so on. They have been overcome. It is certainly a tribute to the VEC that it took on vVote and that it was fully deployed and supported, and offered to very many electors – all of whom where people with special needs regarding voting. This project is the subject of a report published in IEEE Security and Privacy Magazine out this month (July 2016) in the US.

Indirect benefits of verifiable voting

When systems are fully verifiable, layers and layers of security usually needed for secure IT services can actually be relaxed or removed. This may seem odd given we are trying to protect a critical service. However, if we have the ability to detect any and all loss, damage or malfeasance then another powerful effect emerges: there is no hack that can be perpetrated without detection. This means there is no compromise or effect that can take place that will not put the crime and the criminal at risk. Since most crime is crime of opportunity, it should be clear that verification is a great disincentive to tampering. This creates knock---on benefits: developers who know the systems they are writing are going to be able to be deeply verifiable will write better code. It is the logical end---point of developers making tidier code when there are going to be code reviews.

Conclusions

The future for voting in Australia, indeed the solution asked for by both Prime Minister Turnbull and Bill Shorten, is an electronic system than can issue rigorous proof that all is well. It should be able to do this any time, and issue proof to anyone. Ideally the system is open source software as well, since elections must also be open. The system is here, it is vVote, and it can be made ready to serve all Australians.

This submission and others to the Inquiry explain why for a range of reasons, the foreseeable future of safe, decisive high stakes elections does not include Internet voting. If computers are used at all to handle votes in any capacity, then the future must include verifiability, end---to---end. It has been shown that the future can still include automation to make elections faster, more accurate and more accessible with outcomes that are defensible and come with meaningful evidence.

Supporting materials

vVote IEEE Security and Privacy Magazine article draft:
http://arxiv.org/abs/1504.07098

NSW Electoral Commission security report on iVote 2015:
https://www.elections.nsw.gov.au/__data/assets/pdf_file/0019/204058/An_overview_of_the_iVote_2015_voting_system_v4.pdf

Layperson's end---to---end verifiable voting (E2EV) explained:
http://arxiv.org/ftp/arxiv/papers/1504/1504.03778.pdf

vVote open source code project: http://coasca.com this site provides animations of how vVote works

ScanTegrity II https://people.csail.mit.edu/rivest/pubs/CCCEx08.pdf

.....