

A solution for Voter Verifiable e-voting in Ireland

Verification by voter verified paper ballot

Michael McMahon.
draft 0.6
11/7/2005

This note outlines a possible solution to the e-voting problem which describes a pragmatic, fairly simple enhancement to the existing Nedap/Powervote system which should satisfy the demand for a voter-verified-audit-trail (VVAT).

This note is structured with an overview, followed by a more detailed description, and finally, a list of common questions that are often raised with respect to VVAT solutions, are answered.

Overview

The essence of this proposal is that voters are given a printed copy of their electronic vote, which they verify before casting their electronic vote, and which they drop into a traditional ballot box. At the count centre an audit takes place in two stages. First, a random sample of paper ballots is taken from the ballot boxes and compared with the aggregated file of mixed and numbered electronic votes. Each electronic vote, and corresponding paper ballot have a unique ID which allows this comparison to be done. So long as all ballots selected for audit are the same on paper as in the electronic file, then this first audit stage is successful. The second stage of the audit, is to take the officially “certified” (as a result of the first stage of the audit) file of electronic votes and give it to an official observer, who verifies the official results by checking the count with an independent counting system. After the election, any interested party can check the count by using their own counting system to do the same. Figure 1. shows an overview of the processes and the sequencing of each step. Steps are indicated in a chronological order, and it is essential that each step be completed in that order. For example, it is critical that the electronic signature for the electronic vote file be published before the audit commences. Similarly, the random sample must be chosen before any of the ballot modules are loaded into the counting system.

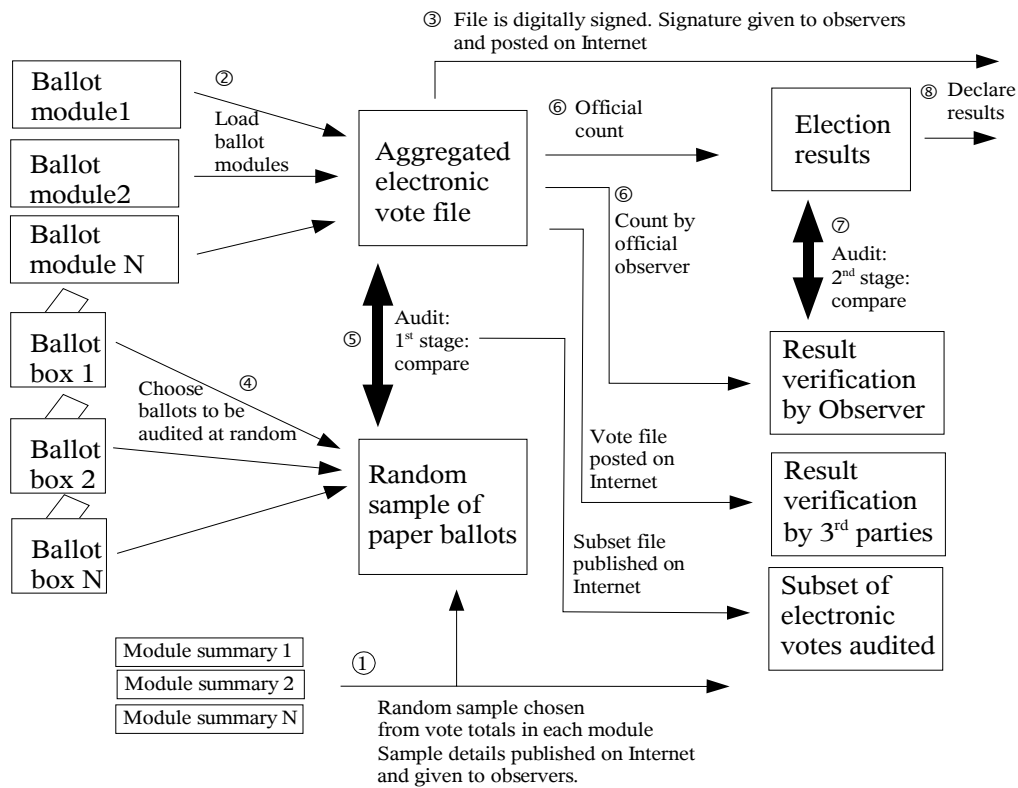


Figure 1 Overview of audit

Description of voter verified audit trail solution

In this solution, a printer is added to the voting machines¹, which prints on paper, a copy of the voters choices as each selection is made. Before printing any vote details, the printer prints out an identification mark at the top of the ballot. This would be equivalent to the mark stamped on old style ballot papers, which authenticates the paper as having been issued by the presiding officer at a particular polling station. As the voter makes each choice, a line is printed showing the candidate and the preference number chosen. If the voter changes his/her mind, then the existing choices are voided by printing a void line across the ballot. The new selections are then printed below the void line and this process can be repeated as necessary until the voter is satisfied that the choices on the machine and on the paper, both represent his/her intended choices. At this point, the “cast vote” button is pressed. Doing this adds a confirmation code and bar-code line to the ballot, and cuts the sheet so the paper can be removed from the printer. The voter must then take the ballot and drop it into a conventional ballot box.

The barcode on the ballot (and the numerical code) encode the voters choices digitally as an aid for the possible auditing of this ballot. In addition to the voters selections, a unique random ballot ID number is included. The Ballot IDs are generated at election set-up time and are unique across the entire set of ballot modules used in a constituency. When the ballot modules are programmed before the election, they would now receive a set of random ballot IDs, and during the election as each voter casts his/her vote, one of these unique IDs is assigned to the voters vote.

¹ the existing voting machine printer is probably not suitable.

These details are important for the selection of ballots to be audited. Some have a difficulty with the notion of recording serial numbers on the ballot paper, due to the risk of linking a particular serial number with the voter who made the vote. The fact is however, that serial numbers are already used on ballot papers in this country and if necessary, the serial numbers can be chosen to be sufficiently large so that it is impractical to record them outside of the electronic system by any of the polling staff, assuming they can surreptitiously see them. In fact, there is a more fundamental, *theoretical* problem for all electronic voting systems than the problem posed by serial numbers. This is the possibility that an electronic voting machine could (secretly) record ballots in the exact sequence in which they are cast (perhaps in addition to the expected random sequence). In theory, the election authorities could use this information to figure out who cast which votes. Ultimately, one has to place *some* trust in the election officials, and one can think of equally likely (or far-fetched, depending on your perspective) scenarios involving traditional paper ballot elections. This could include: checking finger-prints on ballot papers, or possibly installing secret cameras in the polling stations. With electronic voting, one has to accept that while an end-to-end audit trail can prove the accuracy of an election, it can never provide a 100% guarantee of secrecy.

The following simple changes are required in the voting machines:

1. Software changes required to print out the ballots. If external printers are used then hardware changes should not be required, assuming the voting machines have a suitable port, to which a printer can be connected.
2. Software change to the data stored on the ballot modules. The only change here is the addition of the Ballot ID to the data stored with each ballot.

The printer itself can be a fairly simple super-market or Lotto style thermal printer, of the type that has been used reliably for many years now. The ballot box can be of the traditional kind. In fact, the modifications to the existing Nedap machines are minimal.

The system which is used to set-up the election must now be able to generate the ballot IDs and randomly assign them to the complete set of ballot modules which will (or could be) used in the election. Obviously, sufficient numbers of IDs need to be generated to cover all contingencies.

At close of polling

When polling closes, the ballot modules and corresponding ballot boxes are sealed and sent to the count centre. First, the printed statement associated with each ballot module is examined and the total number of votes in each module is noted. A random sample of ballots are “chosen” to be audited as follows. All the ballot boxes are emptied into a large container, the votes are mixed up and a random selection is removed for checking. Next, the ballot modules are loaded into the aggregation/counting system and the official file of mixed and numbered votes is generated. This file must be digitally signed, and the digital signature² provided to all observers (electronically and printed out on paper). After this has been done, the audit comparison can commence. It is important that all present at the count, and the

2 Publishing or distributing the digital signature (before the sample ballots are taken from their boxes) proves that the file (to be audited) has not been modified, in such a way as to bias the audit. If this were not done, then there might be a possibility of altering votes in the electronic file, which were not selected for audit. The Observers will be able to check afterwards that the printed signature is correct (with respect to the published electronic vote file).

observers in particular, can see the sequencing of these events, and in particular can observe the selection of the ballot papers from the container. If the procedure is not visible and transparent, then there could be accusations of fixing the audit sample.

Detailed audit procedure: stage 1

Checking is done with the assistance of a simple PC computer system which includes a keyboard, screen and small hand-held barcode reader. The job can be split among a number of operators (each with own PC), all working in parallel with bundles of ballots divided up amongst them. The operator picks up a ballot and scans the barcode. The digital information on the barcode is displayed on the screen. This consists of the Ballot ID and the voters choices. The operator visually compares the printed data on the ballot with the information displayed on the screen. If the information is the same, then the data is stored in a file which will be examined again later. If the information is not the same, then this indicates a problem with the voting system. This part of the audit can be easily observed by election agents, candidates or other designated observers. In addition to storing the data electronically in a file, a printer is used to simultaneously print out the Ballot ID and vote contents. It is important that the observers be able to see at the same time, the data on the screen, the print-out, and the ballot being audited. At the end of this step of the audit, photo-copies of the printout are given to the observers so they can check the audit later if they wish and the new electronic file containing a subset of votes from the main file is used in the next step.

Next, the subset file of votes is compared electronically (using the Ballot IDs) with the main aggregated vote file. If there are *any* discrepancies then this is evidence of a problem with the voting system. The observers can do this themselves at any later time using their own computer equipment.

Detailed audit procedure: stage 2

The result of stage 1 is effectively the certification of the electronic vote file. This file is then used in an independent implementation of the counting rules in order to verify the official result of the count. If an independent implementation produces the same result as the official system, then there is not likely to be a problem with the counting system.

In addition, the file is published on the Internet so that other interested persons can check the results themselves. The subset vote file, generated by stage 1 of the audit is also published on the Internet. This allows the Observers to compare their paper copy of the file which they received at the audit with the electronic version. The electronic (subset) can then be compared with the complete vote file. This will be done away from the count centre after the election, since it uses data that has to be posted on the Internet by the returning officer.

Note, that no audit errors will occur if the voting system is working correctly. Conversely, if an error is shown, then a manual count would be the only way to continue with the election. Note also, that close results are *not* a reason for automatically triggering recounts, only errors in the audit. Therefore, in all probability, manual counts will *never* need to be done.

Q&A

How many ballots need to be audited and how long would it take?

Given an election with (say) 50,000 ballots, and if you want to be 95% certain of catching errors in 50 or more ballots, then a sample of 3,000 (6%) is sufficient³. Say five operators, each with their own PC auditing system are employed to do the checking, and it takes around five seconds to process each ballot, this would give each operator 600 ballots to process, and this could be done in around 50 minutes. (5 seconds per ballot is realistic because no manual data entry is required).

Isn't this a dual-system with two possible results (one electronic, and one paper)?

This is a common misconception about VVAT. The important thing to understand is that the paper ballot and the electronic vote are both generated by the same software/hardware, and if they are working correctly, then there should be no discrepancy between them. Errors may occur in the audit. For example, an audit operator may mistakenly identify a discrepancy that does not really exist. In this case, a simple recheck of the particular audited ballot will correct the mistake. Alternatively, the operator may fail to notice a discrepancy. In this case, the system depends on one of the observers noticing the mistake. Otherwise, the error goes unnoticed. However, the important thing is that when a system is verifiable, this creates an overwhelming incentive for the system suppliers to get it right.

What is the actual basis for trusting this system?

The basis for trusting the system is that an audit is carried out for every election race using simple procedures that are transparent and observable. There is no need to trust the audit system because it can be independently verified by any of the observers who observe the audit.

What happens if some paper votes are lost?

This is an important issue. If it happens that some ballot boxes are lost or mislaid, then the votes corresponding to these boxes would have to be removed from the count even if the electronic ballots were still present. This may seem a strange course of action, but it would help to underline the position of the paper votes as the actual voting record, and it therefore follows, that no electronic votes can be included in the count unless they can be audited. In practise this situation would be unlikely to occur, since the ballot boxes containing the paper votes would be transported and handled at the same time and in the same way as the electronic ballot modules.

What if a voter claims the machine display and paper ballot are not the same?

There are a number of scenarios where a voter could claim (perhaps mischievously) that the voting machine is not working. This presents a potential problem because election workers would not be allowed to look at the vote to confirm such a problem. A simple procedure is possible for dealing with this. Basically, each voter would be given a time-limit, within which the vote must be cast. If the time expires (for whatever reason) the vote is cancelled and the voter directed back to the presiding officer, who would provide a traditional paper ballot to fill-out. This ballot would then be dropped into a separate ballot box, which would be handled at the count centre as follows. In the presence of multiple witnesses, each such manual ballot would be entered into a voting machine. Obviously, it would have to be made clear to voters that this is *not* an alternative to electronic voting, but just a form of

³ Exact probability of failing to detect at least one error P is $((N-F)!(N-T)!)/(N! (N-F-T)!)$ where F is number of errors, T is number of checks, N is number of ballots.

anonymous assistance which helps people who have difficulty in using it, and which also happens to deal with mischievous voters.

What if a voter fails to post the paper ballot after casting his/her vote electronically?

It is important that voters do not leave the polling station with the paper ballots. This is particularly important in this system because the vote could be counted electronically, and the paper ballot could be used as a receipt for vote-selling purposes. This is one reason why some people advocate the use of screens so that the voter does not get to handle the paper ballots, but since the voting machines in Ireland are individually supervised, it may be sufficient to warn voters that taking ballot papers out of the polling station is illegal. If in spite of this, a voter still manages to leave without posting the paper ballot, a procedure may be required where the last vote on a machine can be cancelled due to the illegal action of the voter. If voters are made aware of all these points, then problems should not occur in practice. Note, see description of Mercuri method below for another possible solution to this problem.

What happens if the voter fails to post the paper ballot, but the machine supervisor does not cancel the vote?

If in spite of the precautions described above, a vote still gets through electronically, but no corresponding paper ballot exists, then the audit will not catch this because the audit procedure only checks ballots which have been verified by the voter *and* posted in the ballot box. This situation could also arise if the machine maliciously adds extra ballots to the electronic file, which were not cast by any voters. This situation has to be checked by comparing the number of votes in the electronic file with the records from the polling station which show how many voters were supposed to have used that machine. This is not really any different from the checks in a traditional system against ballot-stuffing. So long as the number of votes in the electronic file is correct, and most voters correctly post their paper ballots, then the audit will pick up modifications to votes in the electronic file⁴

What if voters post fake ballot papers in the ballot box?

If a voter maliciously posts a fake ballot paper, for the purpose of disrupting the audit, then this situation will be detected in the first instance, if the ballot is chosen for audit, and second if it fails the audit test. The way to deal with this, is to accept the possibility that fake papers may turn up in the audit, and to add sufficient authenticating information to the real papers. This would include serial numbers printed on the back of the paper, together with special designs on the paper itself. Additionally, as already described, an authenticating mark would be printed on each ballot. This mark would be unique to each polling station and each election.

Is there another solution to these problems?

Note, the previous problems can also be dealt with by using the so-called Mercuri method, whereby the printer and ballot box are physically separated from the voter by a perspex screen. If the voter cannot touch the paper ballot, then these problems are automatically solved. However, this inevitably leads to a possibly more costly and complicated arrangement of printer and ballot box. Also, to avoid problems where a voter might claim that the machine printed their vote before it was confirmed, it must be possible for the machine supervisor to see the ballot drop into the ballot box, but without being able to see the information printed on it.

⁴ at the confidence level defined by the sample size

How complicated is the auditing system?

The auditing system is extremely simple. In principle, the audit could be done with a standard computer and using no more than a standard word processor program. This is important because it is essential that elements of the audit can be repeated by the observers in their own time and using their own computer equipment. However, a specially designed audit application, which can read barcodes and display the information in a clear manner, will speed up the audit dramatically. This kind of application is probably at the level of complexity of a Computer Science student assignment.

Doesn't publishing the aggregated vote file risk compromising secrecy?

It has been pointed out that if voters can find their own votes in a published file, then this could lead to vote-selling or coercion scenarios, where a voter is given an unusual (and probably unique) combination of lower preferences, as well as a specific bought (or coerced) high preference. If the unique vote turns up in the file, then the voter has done what he/she was told, thus compromising the secrecy of the ballot. The solution to this problem has two facets. First, the aggregated vote file has to be modified, before it is published and audited, so that for each ballot, any preferences that were not used in the count, are removed from the electronic record. Doing this reduces the uniqueness of votes dramatically and makes it much more difficult to locate individual votes in the file. This solves the problem as stated above. The procedure as described in Fig. 1 needs to be modified as follows. Step 2, which was formerly "load ballot modules" is now the following sub-sequence of steps.

1. load ballot modules to produce a provisional vote file
2. run the count on the provisional vote file
3. the provisional vote file is modified to remove all unused preferences from the count and this modified file becomes the "aggregated vote file" to be published and audited. The provisional file is no longer needed and can be deleted.

Now, the electronic count takes place before any of the auditing stages. Also, the observers and the audit operators will *not* do a complete visual comparison of the paper ballot with the relevant electronic record. They can only compare whatever preferences are revealed in the electronic record with the equivalent preferences on the paper ballot. This should in fact be easier and faster than having to compare the entire ballot. But, because not all of the contents of each ballot are being revealed, this places extra importance on the 2nd stage of the audit, where the count is repeated by the Observer (and possibly by any 3rd party later). Now, the audit 2nd stage must additionally verify that all preferences that are actually needed for the count on each and every electronic ballot are in fact revealed. For example, if a candidate is eliminated after the first count, then it would not be allowable that all following preferences are hidden on that candidates ballots. This would immediately be evidence of a problem with the counting system. On the other hand, if a candidate is elected on the first count, then all ballots not selected for transfer will in fact have all further preferences hidden. Because the count proceeds deterministically⁵ from start to finish it is always therefore possible to know with certainty which preferences need to be visible and which may be hidden.

However, in spite of the procedure just described, there is still a (more remote) possibility that the same kind of fraud could be committed using higher preferences for "no-hope" candidates. In this case, the no-hope votes cannot be removed from the ballots because they are actually used in the count. The way to deal with this is to initially analyse the vote file, looking for elected candidates who received an unusually large number of lower preferences. If there appears to be a pattern of votes going to no-hope

5 once the vote files are already mixed and numbered.

candidates, followed by another candidate who is under suspicion of cheating, then two possible courses of action exist:- first, criminal investigation of vote selling and second, possibly deciding not to publish the vote data. Not publishing would be a last resort because the integrity and trust in the system depends on anybody being able to check the results with an independent counting system. Even though this would be an unusual and drastic step, it might be needed as an option and therefore the function of the observers is even more critical. Given this possibility, it might make sense for the audit function to be taken on by the CEV, as a suitably independent body.

How can the quality of the visual audit be maintained (and measured)?

If, in a particular election, there is no fault in the system, then there will be zero errors detected in the audit. This presents a problem in that the operators, who may be expecting no errors, may rush through the audit, without paying enough attention to it. This situation can be ameliorated by introducing a number of erroneous control ballots into the audit. This can be done without having to modify the electronic vote file.

When the election is set-up and the ballot modules are programmed, the set-up system would reserve a number of ballot Ids as control ballots. These ballot Ids would therefore not be assigned to ballot modules for real votes. The fake control ballots would be printed off at set-up time. The ballots themselves would look identical to those used in the count except that the barcode data is different from the printed data (this being the error to be detected). The returning officer would have to place the control ballots into the bundle of ballots for the audit (at random) and also provide to all observers, before the audit commences a statement showing a list of ballot Ids in the control sequence. The same list could also be deposited with some central authority (before the election), so there is no dispute as to the identity of the control ballots.

With this kind of system, the auditors would know that a certain number of errors will always need to be detected and as each error is detected, it is checked off against the list of control ballots by all observers. Obviously, if any auditor fails to detect a control ballot then that would be cause for concern regarding the quality of the check.

Clearly, precautions would be needed to keep the control ballots separate from the real ones, especially in the case of a manual count. However, the number of control ballots would likely be quite small (perhaps fewer than 20) and the central authority would be arbiter in case of dispute over whether a particular pape is a control or real ballot.

Conclusion

This system is a simple VVAT enhancement to the existing electronic system, which uses routine auditing of every election race, to ensure the electronic system is always functioning correctly. Manual counts would not be required, except in cases where an audit has shown that the system has failed.

The limitations in this system are as follows:

1. Administrative overhead. A cost is associated with auditing in this system. The cost is primarily the management and transport of the paper ballots, together with the cost of doing the manual checking. Compared to the overall election cost, I would estimate this is not excessive. However, there may be a pressure to reduce the amount of auditing to an unacceptably low level, if (as is to be hoped) few problems are found with the system.
2. The audit can only be observed by individuals actually present at the count, and quite a considerable amount of tedious effort is required from the observers (if they want to completely verify the audit). Again compared to the existing system, this is not a significant limitation, but as will be shown, it is possible with electronic voting systems to open up the auditing function so that any interested person, wherever they are located, can do it with no physical effort whatsoever.

The second solution is based on the election verification system proposed by Dr. David Chaum. This solution solves the two problems above. In particular, the administrative overhead is minimal for the election authorities. The only additional cost is the cost of printing out receipts given to each voter as they vote. No paper copies are kept by the election officials. Therefore, there are no costs associated with transport or secure manual handling of auditable material. Verification is done partly by the voters themselves in the polling booth, and later after they leave the polling station.

All auditing is done by publishing particular data on the Internet, and allowing absolutely anyone to download this data and perform various tests on it to verify its correctness. The main problem with this system is the novelty of the concepts behind it, and the fact that it may take some time before it can be developed into a real product. It is also questionable whether the existing Nedap voting machines have sufficient computing power to cope with the additional computation involved with this system. A more detailed description of how Chaum's system might be applied to the Irish (STV) election system will be described in a second part to this paper.

-0-