CS 161 Spring 2010 — Computer Security Paxson/Wagner — Discussion 11

1. Attacking Homework Submission

Cynthia is the only one with a key to the CS161 homework box, but she is not in Soda Hall when CS161 homeworks are due at midnight. Matt is in Soda Hall at midnight (sadly), but he doesn't have a key to the homework box.

(a) How can Cynthia and Matt ensure that homeworks turned in after the deadline are not accepted?

Answer: Matt generates a nonce and writes it on a blank sheet of paper. He tells Cynthia the nonce (over some secure channel). When Cynthia collects the homeworks, she knows that the ones on top of the sheet of paper containing the correct nonce have been turned in late.

(b) How could students defeat this defense?

Answer: If a student puts his or her homework in a different box, it is unclear how the staff should handle this homework submission. It could be an honest mistake, and the staff has no idea when this homework submission was turned in. If they accept it, then they may have accepted a late homework, so they should probably also accept any homeworks known to be late (i.e., found above the correct sheet of paper in the correct box).

2. Election Audits

Naiveville, CA would like to use a post-election audit to provide confidence in the election results. Voters in Naiveville vote on paper ballots which are then read by optical scan machines. The optical scanners produce a tally of votes for each contest. After the election, officials will manually recount the ballots from some percentage of machines. In order to make the audit transparent, the officials publish PDFs of all the voted ballots from each machine as well as the results returned by that machine.

(a) How does Naiveville's policy allow a violation of the secret ballot property? Before the election every ballot for a particular precinct is identical. That is, there are no sequence numbers, or other identifying marks on an un-voted ballot. The voted ballots are shuffled so that the PDFs do not reveal any information about the order in which the votes were placed.

Hint: In CA a single ballot often has many contests on it. Also, many contests allow write-in candidates.

Answer: A voter can sell her vote in the following way. She votes for the candidate buying her vote, and then "signs" her ballot by voting in an unusual pattern on the down-ballot contests. The same technique can also be used to coerce a voter into voting a particular way.

The voter could also try to make some identifying mark in the margin or some other unused portion of the ballot. However, election rules typically specify that any ballot that appears to have been marked

in this way must be rejected. Depending on how strict the rules are, a voter may be able to get away with a small, inconspicuous mark.

(b) How many bits of information can be encoded on a ballot with 10 yes/no contests?

Answer: For each yes/no contest a voter has four options: yes, no, blank, yes & no. So each contest can encode 2 bits of information. Note that usually the optical scanner is programmed to not accept ballots with over-votes (both yes and no selected), but they will usually accept ballots with under-votes (no selection made for a particular contest). In that case, each contest can be used to encode $\frac{\log 3}{\log 2} \approx 1.585$ bits of information. So a ballot with 10 yes/no contests can encode 15-20 bits of information.

(c) Can you think of a non-cryptographic way that election officials can publish the voted ballots while still providing secret ballots?

Answer: After the election, tear each ballot so that every contest is on a separate sheet of paper. Shuffle all the papers from a particular precinct or machine and then publish. A precinct-based or machine-based audit can still be conducted, but any "signature" provided by voters is lost. This does not prevent signatures consisting solely of unusual write-in names or any other identifying marks that are sufficiently subtle to escape detection.

3. Designing BART Cards

(a) The BART system uses a scheme similar to the following. Each card has a magnetic stripe that is read from and written to by the turnstiles. Encoded in the stripe in plain text is an ID and the current dollar amount on the card. To prevent counterfeiting, the card also stores a MAC of the current dollar amount. The readers in the turnstile won't accept the card unless the MAC is valid and and there is enough money on the card. Describe at least one specific attack on the cards.

Answer: Replay attack. Buy a new $100 BART card. Use a magnetic-stripe reader to read the information on the card. Buy some blank cards with magnetic stripes and copy the information to the new cards.

(b) How would you modify the above protocol so it isn't susceptible to the attack you came up with above? You can assume the turnstiles are networked and can contact a central server. What could you do if only a fraction of the turnstiles are networked? What if none of the turnstiles are networked?

Answer: If the turnstiles are networked, there could be a central database that stores a table of card IDs with their current amount. The communication between the turnstiles and the server would also need to be authenticated and would need to include a nonce to prevent replay attacks.

4. Detecting Spam

(a) Many spam filters rely on users to report spam. Once a user has reported an email as spam, the filter tries to learn how to recognize future instances of similar spam. One simple scheme is to remember the from addresses of the spam and not allow any future emails from those addresses. How can the spammer evade this filter?

Answer: Spoof the from address. Use a botnet to send the spam so it is coming from millions of different addresses.

(b) You've probably received spam that has paragraphs of seemingly random words or paragraphs of text taken from works of literature. Why does the spammer do this?

Answer: Spam filters don't rely on the from address to recognize spam. Instead, when user reports an email as spam, the filter tries to learn something about what the message text of spam looks like. By including many random English words or full paragraphs of legitimate text, the spammers are trying to confuse the filters. The point is to force the filters into having a very high false positive rate making them unusable for most people.