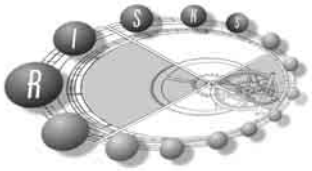

Insider Risks in Elections

Many discussions of voting systems and their relative integrity have been primarily technical, focusing on the difficulty of attacks and defenses. This is only half of the equation: it's not enough to know how much it might cost to rig an election by attacking voting systems; we also need to know how much it would be worth to do so. Our illustrative example uses the most recent available U.S. data, but is not intended to be specific to any particular political parties.

In order to gain a clear majority of the House in 2002, Democrats would have needed to win 13 seats that went to Republicans. According to Associated Press voting data, Democrats could have added 13 seats by swinging 49,469 votes. This corresponds to changing just over 1% of the 4,310,198 votes in these races and under 1/1000 of the 70 million votes cast in contested House races. The Senate was even closer: switching 20,703 votes in Missouri and New Hampshire would have provided Democrats with the necessary two seats.

Of course, it isn't possible to anticipate exactly how much fraud or undetected error would alter the winner of each race. It would also be suspicious if Democrats won 13 districts by exactly one vote. As a result, modestly more votes would need to be changed. In 2002, fraud that changed 2% of the votes in a few contested races (or 1/250 of the total votes) would have completely changed the balance of power in Congress.

According to the Federal Election Commission, some House candidates spent up to $8 million in 2002, although expenditures of $3 to $4 million were typical. Thus, it is easily worth $3 million for a candidate to change a race from a statistical dead heat into a certain victory. Each 1% that is added to a candidate's odds of victory (and hence each 1% removed from the opponent's odds) is worth $60,000.

The outcomes of the 13 closest Democratic losses in 2002 would have changed by swinging an average of 3,805 votes each. If shifting 5,000 votes is worth $3 million, each vote is worth $600. A discount is required to reflect the additional legal risks and moral problems involved in committing fraud, although these effects depend on the people and situations involved. The following analysis makes the conservative assumption of $400 per vote.

So, what is it worth to compromise a voting machine? Suppose one machine collects 250 votes, with roughly half for each candidate in a close election. Rigging the machine to swing all of its votes in one race would be worth $50,000. To avoid detection, fraudsters may be less greedy. Swinging 10% of the opposition's votes on any given machine would be worth $5,000 in a close race. Thus, it is necessary to assume that attacks against individual voting machines are a serious risk, particularly if a few dozen machines could be affected. For example, machine tampering is worthwhile if machines are stored without strong physical security.

Election data is also useful for understanding the threats against voting machine designs. Any voting machine type deployed in 25% of precincts would register enough votes that malicious software could swing the balance of power without creating detectable statistical abnormalities. According to the FEC, Congressional candidates together legally raised over $600 million in 2002. One might conservatively estimate that stealing control over the House of Representatives is worth over $100 million to the party that would otherwise lose. In practice, the threats are even greater, since one attack could affect many elections.

Who are the adversaries? Elections face threats from system developers, election insiders, foreign governments, radical extremists, partisan operatives, and others. Voting systems must be able to face attackers with extraordinary creativity and dedication—much more so than the rather simplistic and unmotivated creators of viruses and worms—because there are strong rational (though perverse) motives for election fraud. Compared with violence and other illegal activities extremists use, electoral fraud is much safer and much more likely to have a desired effect.

The evidence clearly shows voting systems must be designed to counter very well-funded and sophisticated opponents, including those with massive financial resources and the ability to join design teams, infiltrate manufacturing facilities, fabricate malicious integrated circuits, tamper with compilers, and mount a wide range of other attacks. Checks and balances, such as local party observers, help against some attacks but not others. The threats are real, making openness and verifiability critical to election security. C

Paul Kocher (paul@cryptography.com) heads Cryptography Research, Inc.
Bruce Schneier (www.schneier.com) is CTO of Counterpane.

PAUL WATSON