symantec.
press

Crimeware

Understanding New Attacks and Defenses

"The biggest development in online security in the last five years has been the emergence of a criminal economy where villains specialize and trade with each other. This book provides a much-needed update on the tools that these gangsters use now and on others that they might be using in the near future."

—Ross Anderson, Professor of Security Engineering, Cambridge University

Markus Jakobsson • Zulfikar Ramzan

Library of Congress Cataloging-in-Publication Data
Jakobsson, Markus.
Crimeware : understanding new attacks and defenses / Markus Jakobsson, Zulfikar Ramzan. p. cm.
Includes bibliographical references and index.
ISBN 978-0-321-50195-0 (pbk. : alk. paper) 1. Computer security. 2. Internet—Security measures. 3. Computer crimes. I. Ramzan, Zulfikar. II. Title.

QA76.9.A25J325 2008
005.8–dc22 2007050736

ISBN-13: 978-0-321-50195-0
ISBN-10: 0-321-50195-0

Text printed in the United States on recycled paper at Courier in Stoughton, Massachusetts.
First printing, April 2008

Preface

Traditionally, malware has been thought of as a purely technical threat, relying principally on technical vulnerabilities for infection. Its authors were motivated by intellectual curiosity, and sometimes by competition with other malware authors.

This book draws attention to the fact that this is all history. Infection vectors of today take advantage of social context, employ deceit, and may use data-mining techniques to tailor attacks to the intended victims. Their goal is profit or political power. Malware become *crimeware*. That is, malware has moved out of basements and college dorms, and is now a tool firmly placed in the hands of organized crime, terror organizations, and aggressive governments. This transformation comes at a time when society increasingly has come to depend on the Internet for its structure and stability, and it raises a worrisome question: *What will happen next?* This book tries to answer that question by a careful exposition of what crimeware is, how it behaves, and what trends are evident.

The book is written for readers from a wide array of backgrounds. Most sections and chapters start out describing a given angle from a bird's-eye view, using language that makes the subject approachable to readers without deep technical knowledge. The chapters and sections then delve into more detail, often concluding with a degree of technical detail that may be of interest only to security researchers. It is up to you to decide when you understand enough of a given issue and are ready to turn to another chapter.

Recognizing that today's professionals are often pressed for time, this book is written so that each chapter is relatively self-contained. Rather than having each chapter be sequentially dependent on preceding chapters, you can safely peruse a specific chapter of interest and skip back and forth as desired. Each chapter was contributed by a different set of authors, each of whom provides a different voice and unique perspective on the issue of crimeware.

This book is meant for anyone with an interest in crimeware, computer security, and eventually, the survivability of the Internet. It is not meant only for people with a technical background. Rather, it is also appropriate for makers of laws and policies, user interface designers, and companies concerned with user education. The book is not intended as a guide to securing one's system, but rather as a guide to determining what the problem really is and what it will become.

Although we often use recent examples of attacks to highlight and explain issues of interest, focus here is on the underlying trends, principles, and techniques. When the next wave of attacks appears—undoubtedly using new technical vulnerabilities and new psychological twists—then the same principles will still hold. Thus, this book is meant to remain a useful reference for years to come, in a field characterized by change. We are proud to say that we think we have achieved this contradictory balance, and we hope that you will agree.

Acknowledgments

We are indebted to our expert contributors, who have helped make this book what it is by offering their valuable and unique insights, and selflessly donated their time to advance the public's knowledge of crimeware. The following researchers helped us provide their view of the problem: Shane Balfe, Jeffrey Bardzell, Shaowen Bardzell, Dan Boneh, Fred H. Cate, David Cole, Vittoria Colizza, Bruno Crispo, Neil Daswani, Aaron Emigh, Peter Ferrie, Oliver Friedrichs, Eimear Gallery, Mona Gandhi, Kourosh Gharachorloo, Shuman Ghosemajumder, Minaxi Gupta, James Hoagland, Hao Hu, Andrew Kalafut, Gary McGraw, Chris J. Mitchell, John Mitchell, Steven Myers, Chris Mysen, Tyler Pace, Kenneth G. Paterson, Prashant Pathak, Vinay Rao, Jacob Ratkiewicz, Melanie Rieback, Sourabh Satish, Sukamol Srikwan, Sid Stamm, Andrew Tanenbaum, Alex Tsow, Alessandro Vespignani, Xiaofeng Wang, Stephen Weis, Susanne Wetzel, Ollie Whitehouse, Liu Yang, and the Google Ad Traffic Quality Team.

In addition, Markus wishes to thank his graduate students, who have helped with everything from performing LaTeX conversions to being experiment subjects, and many of whose research results are part of this book. Zulfikar wishes to thank Oliver Friedrichs and the rest of the Symantec Advanced Threat Research team (as well as his colleagues throughout Symantec) for affording him the opportunity to work on this book and for engaging in countless stimulating discussions on these topics.

We also both want to acknowledge the help and guidance we have received from Jessica Goldstein and Romny French at Addison-Wesley.

Finally, we want to thank our understanding spouses and families, who have seen much too little of us in the hectic months during which we labored on getting the book ready for publication.

Markus Jakobsson
Palo Alto, California
January, 2008

Zulfikar Ramzan
Mountain View, California
January, 2008

Chapter 10

Cybercrime and Politics

Oliver Friedrichs

While we first saw the Internet used extensively during the 2004 U.S. presidential election, its use in future presidential elections will clearly overshadow those humble beginnings. It is important to understand the associated risks as political candidates increasingly turn to the Internet in an effort to more effectively communicate their positions, rally supporters, and seek to sway critics. These risks include, among others, the dissemination of misinformation, fraud, phishing, malicious code, and the invasion of privacy. Some of these attacks, including those involving the diversion of online campaign donations, have the potential to threaten voters' faith in the U.S. electoral system.

The analysis in this chapter focuses on the 2008 presidential election to demonstrate the risks involved, but our findings may just as well apply to any future election. Many of the same risks that we have grown accustomed to on the Internet can also manifest themselves when the Internet is expanded to the election process.

It is not difficult for one to conceive of numerous attacks that might present themselves and, to varying degrees, influence the election process. One need merely examine the attack vectors that already affect consumers and enterprises today to envision how they might be applied to this process. In this chapter, we have chosen to analyze those attack vectors that would be most likely to have an immediate and material effect on an election, affecting voters, candidates, or campaign officials.

A number of past studies have discussed a broad spectrum of election fraud possibilities, such as the casting of fraudulent votes [258] and the security, risks, and challenges of electronic voting [173]. There are many serious and important risks to consider related both to the security of the voting process and to the new breed of electronic voting machines that have been documented by others [46]. Risks include the ability for attackers or insiders either to manipulate these machines or to alter and tamper with the end results. These concerns apply not only to electronic voting in the United States, but have also been raised by other countries, such as the United Kingdom, which is also investigating and raising similar concerns surrounding electronic voting [274]. Rather than revisit the subject of electronic voting, the discussion here focuses exclusively on Internet-borne threats, including how they have the potential to influence the election process leading up to voting day.

We first discuss domain name abuse, including typo squatting and domain speculation as it relates to candidate Internet domains. Next, we explore the potential impact of phishing on an election. We then discuss the effects of security risks and malicious code, and the potential for misinformation that may present itself using any of these vectors. Finally, we review how phishers may spoof political emails (such as false campaign contribution requests) instead of emails appearing to come from financial institutions. The goal in such attacks might still be to collect payment credentials, in which case the political aspect is just a new guise for fraud. However, political phishing emails might also be used to sow fear among potential contributors and make them less willing to contribute online—whether to spoofed campaigns or to real ones.

These sets of risks cross technical, social, and psychological boundaries. Although traditional forms of malicious code certainly play an important role in these threats, social engineering and deception provide equal potential to be exploited and might have a more ominous psychological impact on voters who are exercising their right to elect their next president, or cast their vote in any other type of election.

This chapter includes both active research conducted by the author and discussion of how current threats may be customized. To determine the impact of typo squatting and domain name speculation, for example, we performed an analysis of 2008 presidential election candidate web sites and discovered numerous examples of abuse.

In regard to the attacks discussed in this chapter, we believe and hope that candidates and their campaigns are unlikely to knowingly participate in or support these activities themselves, for two reasons. First, it would not be acting in good faith. Second, their actions would in many cases be considered a breach of either existing computer crime or federal election law.[1]

We conclude that perpetrators would likely fall into two categories: those with political motives and those seeking to profit from these attacks. In the end, it may

1. U.S. Code Title 18, Part I, Chapter 29. Available from `http://www4.law.cornell.edu/uscode/html/uscode18/usc_sec_18_00000594----000-.html`

be difficult to identify from a given attack which one of these goals is the attacker's true motive.

10.1 Domain Name Abuse

To communicate with constituents and supporters, candidates have created and maintain web sites, which are identified by and navigated to via their registered domain names. All candidates for the 2008 federal election have registered, or already own, unique domain names that are used to host their respective web sites. In all cases this domain name incorporates their own name in some capacity, and in some cases has been registered specifically in support of the 2008 campaign. Domain names play one of the most important roles in accessing a web site. They are the core part of the URL that is recognized by the general population and, as such, their ownership dictates who can display content to users visiting web sites hosted on that domain name.

While users may well know the URL for their bank or favorite commerce site, voters may not readily know the URL for their political party's or chosen candidate's web site. Legitimate-sounding domain names may not be as they appear. The authors of this book, for example, were able to freely register domain names such as `http://www.democratic-party.us` and `http://www.support-gop.org` that have for some time warned visitors about the risks presented by phishing. It would be easy to use a domain name of this type for the purposes of phishing or crimeware installation.

Consider, for example, an email pointing to one of these domains that contains text suggesting it came from the Democratic Party and asking the recipient for a donation. If willing to contribute, the recipient may be offered to choose a variety of payment methods, each one of which would allow the phisher to potentially capture the user's credentials as he or she enters this data on the site (or on another, suitably named site hyperlinked from the donation page). The email might also offer the recipient a chance to download and access resources, such as campaign movies, which themselves might contain malware. Existing movies can be modified to incorporate malware [388]. Typical Internet users are also very susceptible to attacks in which self-signed certificates vouch for the security of executables as long as a person known to them has also indicated that the material is safe [388]. In one study [388], that known person was a friend; in our hypothetical case, it might be a political party or a politician.

In today's online environment, individuals and businesses must consider a number of risks posed by individuals attempting to abuse the domain name system. These involve domain speculators, bulk domain name parkers, and typo squatters.

10.1.1 Background

Since the early days of Internet commerce, Internet domain names have held an intrinsic value, much as real estate in the physical world has been valued for centuries. In the early 1990s, when relatively few `.com` domain names existed, it was highly probable that if one attempted to acquire the name of a well-known company, individual, or trademark, this name would be readily available. Many early domain name speculators did, in fact, acquire such domain names, in many cases later selling them to the legitimate trademark holder. At that point, the legal precedence for domain name disputes had not yet been set, and the speculator had a chance of profiting from this sale, in particular if it was to a well-known and well-funded corporation.

It was only a matter of time before formal dispute guidelines were created to eliminate such infringement. A formal policy was created by ICANN in 1999, which is known as the Uniform Domain Name Dispute Resolution Policy (UDRP) [127]. The UDRP is implemented in practice by the World Intellectual Property Organization's (WIPO) Arbitration and Mediation Center.

While this policy provides a framework for resolving infringement, it does not preclude the registration of an infringing domain name if that domain name is unregistered. What is in place is a policy and framework for the legitimate trademark owner to become the owner of the domain, granted the trademark owner first becomes aware of the infringing domain's existence. The policy is frequently used by legitimate business trademark holders to protect their names.[2]

While it is used to protect trademarked proper names, the same policy applies to unregistered, or "common law" marks, including well-known individuals' proper names, even when a formal trademark does not exist. Julia Roberts, for example, was able to obtain ownership of the `juliaroberts.com` domain name, even in the absence of a registered trademark.[3] This is common when a domain name is specific enough and matches a full proper name. In other examples, such as the more general domain name `sting.com`, contested by the well-known singer Sting, the transfer was not granted and the original registrant retained ownership.[4]

There appear to be very few cases in which either elected or hopeful political candidates have disputed the ownership of an infringing domain name. One

2. *The Coca-Cola Company v. Spider Webs Ltd.* `http://www.arb-forum.com/domains/decisions/102459.htm`.

3. *Julia Fiona Roberts v. Russell Boyd.* `http://www.wipo.int/amc/en/domains/decisions/html/2000/d2000-0210.html`.

4. *Gordon Sumner, p/k/a Sting v. Michael Urvan.* `http://www.wipo.int/amc/en/domains/decisions/html/2000/d2000-0596.html`

example that does exist is for the domain name `kennedytownsend.com` and several variations thereof. Disputed by Kathleen Kennedy Townsend, who was Lieutenant Governor of the State of Maryland at the time, the transfer was not granted, based predominantly on what appears to be a technicality of how the dispute was submitted. Central to the ruling in such dispute cases is whether the trademark or name is used to conduct commercial activity, and thus whether the infringement negatively affects the legitimate owner and, as a result, consumers:

> Here, the claim for the domain names is brought by the individual politician, and not by the political action committee actively engaged in the raising of funds and promotion of Complainant's possible campaign. Had the claim been brought in the name of the Friends of Kathleen Kennedy Townsend, the result might well have been different. But it was not. The Panel finds that the protection of an individual politician's name, no matter how famous, is outside the scope of the Policy since it is not connected with commercial exploitation as set out in the Second WIPO Report.[5]

Within the United States, trademark owners and individuals are further protected by the Anticybersquatting Consumer Protection Act, which took effect on November 29, 1999.[6] The ACPA provides a legal remedy by which the legitimate trademark owner can seek monetary damages in addition to the domain name, whereas the UDRP provides for only recovery of the domain name itself.

Even today, the relatively low cost involved in registering a domain name (less than $10 per year) continues to provide an opportunity for an individual to profit by acquiring and selling domain names. The relative scarcity of simple, recognizable "core" domain names has resulted in the development of a significant after-market for those domain names and led to the creation of a substantial amount of wealth for some speculators [377]. Today, a number of online sites and auctions exist explicitly to facilitate the resale of domain names.

In addition to engaging in domain name speculation for the purpose of its future sale, many speculators seek to benefit from advertising revenue that can be garnered during their ownership of the domain name. These individuals—and, more recently, for-profit companies such as iREIT[7]—may register, acquire, and own hundreds of thousands to millions of domain names explicitly for this purpose. These domains display advertisements that are, in many cases, related to the domain name itself, and their owners receive an appropriate share of the

5. *Kathleen Kennedy Townsend v. B. G. Birt.* `http://www.wipo.int/amc/en/domains/decisions/html/2002/d2002-0030.html`

6. Anticybersquatting Consumer Protection Act. `http://thomas.loc.gov/cgi-bin/query/z?c106:S.1255.IS:=`

7. Internet REIT. `http://www.ireit.com/`

advertising revenue much like any web site participating in CPM, CPC, or CPA[8] advertising campaigns.

10.1.2 Domain Speculation in the 2008 Federal Election

Typo squatting seeks to benefit from a mistake made by the user when entering a URL directly into the web browser's address bar. An errant keystroke can easily result in the user entering a domain name that differs from the one intended. Typo squatters seek to benefit from these mistakes by registering domain names that correspond to common typos. Whereas in the past users making typos were most likely to receive an error indicating that the site could not be found, today they are likely to be directed to a different web site. In many cases, this site may host advertisements, but the potential for more sinister behavior also exists.

To determine the current level of domain name speculation and typo squatting in the 2008 federal election, we performed an analysis of well-known candidate domain names to seek out domain speculators and typo squatters. First, we identified all candidates who had registered financial reports with the Federal Election Commission for the quarter ending March 31, 2007.[9] A total of 19 candidates had submitted such filings. Next, we identified each candidate's primary campaign web site through the use of popular search engines and correlated our findings with additional online resources to confirm their accuracy. This, in turn, gave us the primary registered domain name upon which the candidate's web site is hosted.

To simplify our analysis, we removed domains that were not registered under the `.com` top-level domain. This resulted in the removal of two candidates who had domains registered under the `.us` top-level domain. Our decision to focus on the `.com` top-level domain was driven by no other reason than our ability to access a complete database of `.com` registrants at the time of our research. Our final list of candidate web sites and their resulting domains appears in Table 10.1.

Once we had identified the set of candidate domain names, we conducted two tests to examine current domain name registration data. First, we determined how widespread the behavior of typo squatting was on each candidate's domain. Second, we examined domain name registration data so as to identify cousin domain names [198]. For our search, we defined a cousin domain name as one that contains

8. See Chapter 11 for a description of CPM, CPC, and CPA, along with a discussion of Internet advertising.

9. FEC Filing from Prospective 2008 Presidential Campaigns. `http://query.nictusa.com/pres/2007/Q1`

Joe Biden (Democrat)	`http://www.joebiden.com`
Sam Brownback (Republican)	`http://www.brownback.com`
Hillary Clinton (Democrat)	`http://www.hillaryclinton.com`
John Cox (Republican)	`http://www.cox2008.com`
Christopher Dodd (Democrat)	`http://www.chrisdodd.com`
John Edwards (Democrat)	`http://www.johnedwards.com`
James Gilmore (Republican)	`http://www.gilmoreforpresident.com`
Rudy Giuliani (Republican)	`http://www.joinrudy2008.com`
Mike Huckabee (Republican)	`http://www.mikehuckabee.com`
Duncun Hunter (Republican)	`http://www.gohunter08.com`
John McCain (Republican)	`http://www.johnmccain.com`
Barack Obama (Democrat)	`http://www.barackobama.com`
Ron Paul (Republican)	`http://www.ronpaul2008.com`
Bill Richardson (Democrat)	`http://www.richardsonforpresident.com`
Mitt Romney (Republican)	`http://www.mittromney.com`
Tom Tancredo (Republican)	`http://www.teamtancredo.com`
Tommy Thompson (Republican)	`http://www.tommy2008.com`

Table 10.1: The final candidate web site list, together with the domain names.

the candidate domain name in its entirety, with additional words either prefixed or appended to the candidate domain name. In this context, we would consider domain names such as `presidentbarackobama.com` or `presidentmittromney.com` as cousin domain names to the candidates' core domain names of `barackobama.com` and `mittromney.com`, respectively. One can also define a cousin name more loosely as a name that semantically or psychologically aims at being confused with another domain name. In this sense, `www.thompson-for-president.com` should be considered a cousin name domain of `www.tommy2008.com`, despite the fact that they do not share the same core. For the sake of simplicity, we did not examine cousin domains that are not fully inclusive of the original core domain name.

To generate typo domain names, we created two applications, `typo_gen` and `typo_lookup`. The `typo_gen` application allowed us to generate typo domain names based on five common mistakes that are made when entering a URL into the web browser address bar [466].

Missing the first "." delimiter:	`wwwmittromney.com`
Missing a character in the name ("t"):	`www.mitromney.com`
Hitting a surrounding character ("r"):	`www.mitrromney.com`
Adding an additional character ("t"):	`www.mitttromney.com`
Reversing two characters ("im"):	`www.imttromney.com`

As a result of such mistakes, the potential number of typos grows in proportion to the length of the domain name itself. The sheer number of typos for even a short domain name can be large. It is rare to find that an organization has registered all potential variations of its domain name in an effort to adequately protect itself. Typo squatters take advantage of such omissions to drive additional traffic to their own web properties.

Our second application, `typo_lookup`, accepted a list of domain names as input and then performed two queries to determine whether that domain name has been registered. First, a DNS lookup was performed to determine whether the domain resolves via the Domain Name System (DNS). Second, a `whois` lookup was performed to identify the registered owner of the domain.

For the purposes of our analysis, we considered a domain to be typo squatted if it was registered in bad faith by someone other than the legitimate owner of the primary source domain name. We visited those web sites for which typos currently exist and confirmed that they were, in fact, registered in bad faith. We filtered out those that directed the visitor to the legitimate campaign web site as well as those owned by legitimate entities whose name happens to match the typo domain.

Our second test involved the analysis of domain registration data to identify cousin domain names. We obtained a snapshot of all registered domains in the `.com` top-level domain during the month of June 2007. We performed a simple text search of this data set in an effort to cull out all matching domains.

Additional techniques could be used to generate related domain names that we did not examine during our research. This may include variations on a candidate's name (`christopher` instead of `chris`), variations including only a candidate surname (`clinton2008.com`), and the introduction of hyphens into names (`mitt-romney.com`). In addition, a number of typos might be combined to create even more variations on a given domain name, although it becomes less likely that an end user will visit such a domain name as the number of mistakes increases. Nevertheless, such domain names can be very effective in phishing emails, because the delivery of the malicious information relies on spamming in these cases, and not on misspellings made by users.

Expanding our search criteria in the future may result in the discovery of an even larger number of related domains. It also has the side effect of increasing our false-positive rate, or the discovery of domains that appear related but may, in fact, be legitimate web sites used for other purposes. In addition, the amount of manual analysis required to filter out such false positives further forced us to limit our search. Our results are shown in Table 10.2.

We can draw two clear conclusions from the results of our analysis. First, a large number of both typo and cousin domain names were registered by parties other than the candidate's own campaign. We found that many of the registered

Domain Name	Registered Typo Domains	Example	Registered Cousin Domains	Example
barackobama	52 of 160	narackobama	337	notbarackobama
brownback	0 of 134		152	runagainstbrownback
chrisdodd	14 of 145	chrisdod	21	chrisdoddforpresident
cox2008	3 of 92	fox2008	50	johncox2008
gilmoreforpresident	0 of 276		20	jimgilmore2008
gohunter08	1 of 150	ohunter08	23	stopduncanhunter
hillaryclinton	58 of 191	hillaryclingon	566	blamehillaryclinton
joebiden	15 of 125	jobiden	43	firejoebiden
johnedwards	34 of 170	hohnedwards	190	goawayjohnedwards
johnmccain	20 of 137	jhnmccain	173	nojohnmccain
joinrudy2008	9 of 173	jionrudy2008	123	dontjoinrudy2008
mikehuckabee	3 of 167	mikehukabee	28	whymikehuckabee
mittromney	18 of 123	muttromney	170	donttrustmittromney
richardsonforpresident	2 of 340	richardsonforpresiden	69	nobillrichardson
ronpaul2008	11 of 143	ronpaul20008	276	whynotronpaul
teamtancredo	1 of 170	teamtrancredo	16	whytomtancredo
tommy2008	1 of 107	tommyt2008	30	notommythompson

Table 10.2: Typo squatting and cousin domain analysis results. Many typo domain names were already registered and being used in bad faith. In addition, even more cousin domain names were registered, both in support of a candidate and, in many cases, to detract from a candidate. Note that all domains and examples are in the `.com` top-level domain.

web sites, in both the typo squatting case and the cousin domain name case, were registered for the purpose of driving traffic to advertising web sites.

Second, candidates have not done a good job in protecting themselves by proactively registering typo domains to eliminate potential abuse. In fact, we were able to find only a single typo web site that had been registered by a candidate's campaign: `http://www.mittromny.com`. All typo domains were owned by third parties that appeared unrelated to the candidate's campaign.

One observation that we made is that many of the typo domains that displayed contextual advertisements were, in fact, displaying advertisements that pointed back to a candidate's legitimate campaign web site. This is best demonstrated in Figure 10.1. In such cases, a typo squatter had taken over the misspelling of a candidate's domain name and was able to profit from it. Even worse, the candidate was paying to have his or her ads displayed on the typo squatter's web site! This is a result of the way in which ad syndication on the Internet works.

Figure 10.1: When we visited `http://www.barackobams.com` (a typo of Barack Obama's web site, `http://www.barackobama.com`), it contained advertisements pointing to the candidate's legitimate campaign site.

Ad syndicates display advertisements on a web site by indexing its content and displaying advertisements that are appropriate given that content. They may also look at the domain name itself and display advertisements for matching keywords in the domain name. As a result, advertisements for the legitimate campaign may be displayed on a typo squatter's web site. When a user mistypes the web site name and browses to the typo domain, he or she is presented with an advertisement for the legitimate campaign's web site. If the user clicks on this advertisement, the ad syndicate generates a profit, giving a portion to the typo squatter for generating the click through and charging the advertiser, which in this case is the legitimate campaign.[10]

10. A more detailed discussion of how Internet advertising works can be found in Chapter 11.

Individuals who register cousin domain names may have similar motives to those of typo squatters, but they may also be speculating on the value of the domain name itself, with the intent to resell it at a later date. It is also possible that they intend to use the domain to defraud people or to make people wary of emails purportedly coming from a given candidate.

In our analysis, the majority of the identified domains, both in the typo and cousin cases, likely had been acquired in bulk, for the explicit purpose of driving traffic to advertisements. As a result, many of these domains were parked with companies that provide a framework for domain name owners to profit from the traffic that their web sites receive.

10.1.3 Domain Parking

Typo squatters and domain name speculators need not host the physical web infrastructure required to display their own web content or to host their advertisements. Instead, domain name owners can rely on domain parking companies that will happily handle this task for them, for an appropriate share of the advertising revenue. Domain name parking companies will provide the required web site and leverage their preestablished relationships with advertising providers to make life as simple as possible for domain name owners. To leverage a domain name parker, the domain name owner need only configure his or her domain's primary and secondary DNS servers to that of the domain parker. This makes the acquisition and profit from the ownership of a domain name even simpler, to the extent that an individual need just register a domain name and park it at the same time.

While registering a domain name and parking that domain name put the core requirements and relationships in place for a revenue generation model, they do not guarantee that the domain owner will, in fact, profit from this setup. To generate a profit, an adequate amount of traffic and interest must be generated to draw Internet users to that domain name. As such, more emphasis is placed on domain names that are more likely to generate more interest. This is supported by our analysis in Table 10.1, which clearly demonstrates that typo squatters and speculators have favored the domain names of leading candidates.

10.1.4 Malicious Intent

While advertising has been the primary motive behind the registration of typo and cousin name domains to date, more measurable damage using these techniques is highly likely to occur. We have already observed a number of cases where a

Figure 10.2: `http://www.hillaryclingon.com` is a typo-squatted version of Hillary Clinton's real web site, `http://www.hillaryclinton.com` (the "g" key is right below the "t" key on the keyboard), but it has another meaning as well.

typo-squatted domain has been forwarded to an alternative site with differing political views, as seen in Figures 10.2, 10.3, and 10.4. This is problematic in the typo squatting case, because the end user is unknowingly being redirected to a different web site. It is even more common when analyzing cousin domains, which can be registered by anyone; the number of possible registrations can become nearly infinite. It is, however, much more difficult to drive visitors to those domains without having some way in which to attract them. As such, owners of cousin domains use other techniques to attract visitors, including manipulating search engines to increase their ranking (search engine optimization) or, in some cases, even taking out their own advertisements. It may also involve phishing-style spamming of a large number of users.

One interesting side effect of ad syndication networks as they exist today is that we frequently encounter typo domains that are hosting advertisements for a candidate's competitor. It is interesting to see how search engine optimization and keyword purchasing play roles in attracting visitors. Many search engines allow the purchasing of advertisements that are displayed only when users search for specific keywords. Google AdWords is a popular example of such a program where particular keywords can be purchased and advertisements of the purchaser's

Figure 10.3: `http://www.joinrudy20008.com`, a typo-squatted version of Rudy Giuliani's campaign web site, `http://www.joinrudy2008.com`, redirects users to a detractor's web site at `http://rudy-urbanlegend.com`.

choice will then be displayed. As shown in Figure 10.5, this may result in advertisements for one candidate being displayed when a user is searching for a particular keyword, or accidentally browsing to a typo-squatted web site.

Advertising, misdirection, and detraction aside, the real potential for future abuse of typo and cousin domains may revolve around the distribution and installation of security risks and malicious code. This attack vector is by no means new, as web sites and banner advertisements are frequently used to attack visitors who happen to browse to a malicious web site [233]. Attackers who control such web sites frequently leverage a software vulnerability in the web browser [234], or use social engineering and misleading tactics to trick the user into installing security risks [95] and malicious code. Even in the absence of a software vulnerability, we can conceive of a number of convincing scenarios that an attacker might use to convince visitors to install such software. For example, a site could easily mirror Hillary Clinton's legitimate web site, but prominently feature an offer for a Hillary Clinton screensaver that is, in fact, spyware or malicious code.

Figure 10.4: `http://www.muttromney.com` is a typo-squatted version (the “u” key is beside the “i” key on the keyboard) of Mitt Romney’s web site, `http://www.mittromney.com`, which redirects the user to a detractor’s web site.

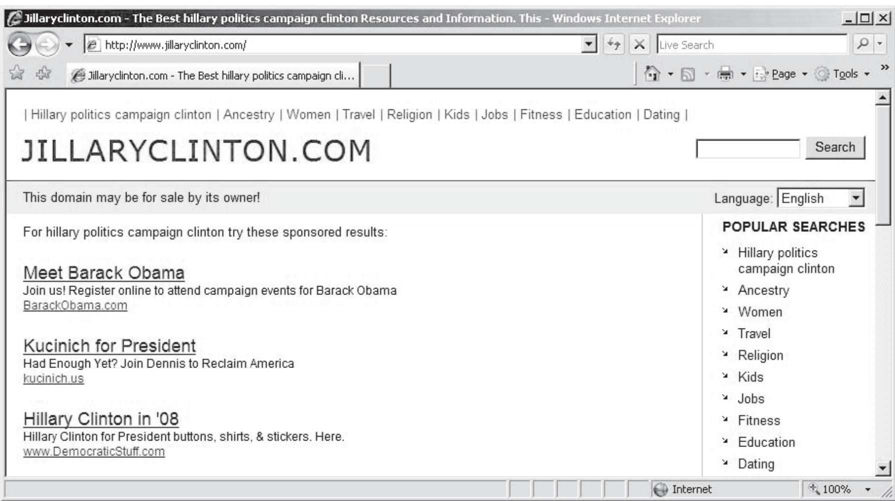

Figure 10.5: `http://www.jillaryclinton.com`, a typo-squatted version of Hillary Clinton’s web site, `http://www.hillaryclinton.com`, displays advertisements directing visitors to rival web sites.

Another site, perhaps mirroring that of Rudy Giuliani, might offer an application claiming to give instant access to his travels, speeches, and videos. Yet another site might claim that by downloading an application, the visitor can assist the candidate in fundraising; that application would, instead, monitor and steal the victim's own banking credentials. The impact of downloading such an application under false pretenses is covered in more detail later in this chapter.

10.2 Campaign-Targeted Phishing

Phishing has without a doubt become one of the most widespread risks affecting Internet users today. When we look at phishing and the role that it may play in an election campaign, we can readily envision several incremental risks that present themselves beyond the traditional theft of confidential information.

10.2.1 Profit-Motivated Phishing

Profit-motivated, event-based phishing is certainly not new. It has been seen in the past on numerous occasions leading up to and following significant events worldwide. For example, this type of attack was seen after natural disasters such as the Indian Ocean tsunami in 2004 [66] and Hurricane Katrina in 2005 [220, 251]. It was also seen in conjunction with sporting events, such as the 2006 and 2010 FIFA World Cup [275].

Election-related phishing has been observed in the past. During the 2004 federal election, phishers targeted the Kerry–Edwards campaign [370], a campaign that was acknowledged as being at the forefront of leveraging the Internet for communications. At least two distinct types of phishing were observed. In one case, phishers set up a fictitious web site to solicit online campaign contributions shortly after the Democratic National Convention; this site stole the victim's credit card number, among other information. In the second case, phishers asked recipients to call a for-fee 1-900 number, for which the victim would subsequently be charged $1.99 per minute [417]. This is a prime example of how such attacks can cross technology boundaries to appear even more convincing. The perpetrators of these two attacks were never caught.

When considering the 2004 election as a whole, phishing presented only a marginal risk. At the time, phishing was still in its infancy, and had yet to grow into the epidemic that can be observed today. When assessing the potential risk of phishing in conjuction with the 2008 federal election, however, we find ourselves in a much different position. Candidates have flocked to the Internet, seeing it as a key means to communicate with constituents and to raise campaign contributions.

Domain Name	Redirects to
barackobama.com	https://donate.barackobama.com
brownback.com	https://www.campaigncontribution.com
chrisdodd.com	https://salsa.wiredforchange.com
cox2008.com	https://www.completecampaigns.com
mikehuckabee.com	https://www.mikehuckabee.com
gilmoreforpresident.com	https://www.gilmoreforpresident.com
gohunter08.com	https://contribute.gohunter08.com
hillaryclinton.com	https://contribute.hillaryclinton.com
joebiden.com	https://secure.ga3.org
johnedwards.com	https://secure.actblue.com
johnmccain.com	https://www.johnmccain.com
joinrudy2008.com	https://www.joinrudy2008.com
mittromney.com	https://www.mittromney.com
richardsonforpresident.com	https://secure.richardsonforpresident.com
ronpaul2008.com	https://www.ronpaul2008.com
teamtancredo.com	https://www.campaigncontribution.com
tommy2008.com	https://secure.yourpatriot.com

Table 10.3: An analysis of 2008 federal candidate web sites and the sites to which contributors are directed to. The sites to which contributors are redirected are legitimate, but the fact that they are often different from the original site increases the risk for confusion and thereby the risk that a phishing attack with a similar design would succeed.

We performed an analysis of campaign web sites in an attempt to determine to what degree they allow contributions to be made online. We discovered that every candidate provided a mechanism by which supporters could make a donation online. All of the web sites on which contributions could be made leveraged SSL as a means to secure the transaction. We also noted the domain of each contribution site. In numerous cases, would-be contributors were redirected to a third-party site, which sat on a different primary domain. Table 10.3 lists both the original domain, and the web site to which the user is redirected.

This redirection was the result of third-party consulting, media, and online advocacy firms being used to assist in the running of the campaign, including the processing of online campaign contributions. This practice does not present a security risk in and of itself, nor is it an indication that phishing is taking place; however, the change in the top-level domain may add to the confusion of potential contributors, who tend to err on the side of caution. It also indicates that additional parties may be involved in the gathering and processing of personal information on behalf of a campaign, increasing the overall exposure of the credit card numbers processed during fundraising.

It should also be noted that the redirection used here is not necessary, and that the contribution site could just as easily remain in the same top-level domain, as a subdomain hosted by the third party for processing. To do so simply requires the appropriate configuration of the primary domain's DNS records. In fact, the majority of the remaining candidates have chosen to follow this path. Future research may also reveal whether those donation sites that do live under the campaign's domain name are, in fact, hosted on the same physical network as the campaign web site or on another third-party payment processor's network.

Figure 10.6 provides a sample of the information collected during an online contribution. We found that forms were fairly consistent in the type of information that was collected, while (not surprisingly) varying from a visual perspective.

The ability to process credit card transactions on an authentic campaign web site may provide an unexpected benefit to online identity thieves. One tactic

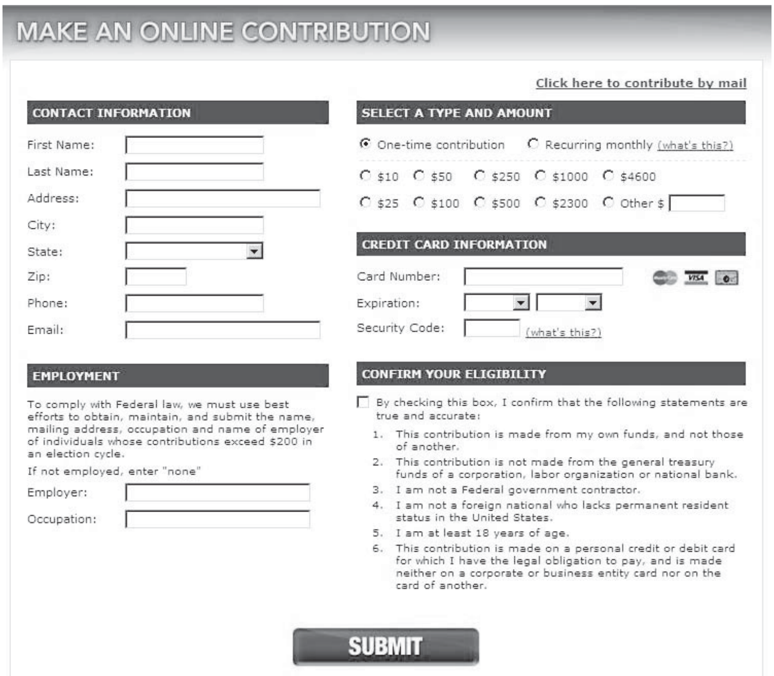

Figure 10.6: A sample form from one candidate's web site allowing visitors to make contributions online. This is a legitimate site. Given that typical Internet users would not be well acquainted with the domains associated with political candidates, there is a risk that phishers might use a similarly designed web site to collect credentials from unsuspecting victims.

regularly employed by those peddling in stolen credit cards is to process a very small transaction so as to validate a credit card as legitimate [48]. Thieves began using this technique in early 2007 on online charity web sites, but it has long been used on other types of online payment sites. Such a small transaction is unlikely to be noticed by the credit card holder and is unlikely to be flagged by the party processing the transaction.

Of course, not all contributions would necessarily be helpful. Attackers might seek to disrupt a candidate's fundraising efforts by initiating illegitimate payments to create confusion. If performed en masse, the widespread contribution of small, random amounts of money, from thousands or tens of thousands of stolen credit cards, would certainly have a negative effect. While there is a slight chance such an attack might remain stealth, it is more likely that it will be noticed, making it nearly impossible to differentiate legitimate contributions from fraudulent donations. Thus a significant burden would be placed on the affected candidates by diluting legitimate contributions with those that were not initiated by the credit card owners.

The increased collection of online campaign contributions also provides a ripe opportunity for phishers to target members of the unsuspecting public. Candidates and their parties regularly communicate with voters through email, as demonstrated in Figure 10.7. Phishing involves the use of email to lure a victim to a fictitious web site that attempts to steal confidential information from the victim [91]. While it is unreasonable to expect campaigns not to solicit contributions using email as a medium, they would be well advised to follow best practices that have been set by other online entities heavily prone to phishing. (A number of excellent resources are available through the Anti-Phishing Working Group [313], including a report funded by the U.S. Department of Homeland Security [101] that discusses the problem in depth and suggests best practices for organizations to communicate safely with their constituents.) However, whether or not the candidate uses email for contribution requests, a phisher may pose as a candidate and ask the recipients of his or her email for money. The typical goal would be to steal the credentials of the victims.

Phishers can increase their success rate by registering domain names that are typos or cousin domains of their target, a tactic already discussed in some depth in this chapter. For example, a phisher targeting John Edwards might elect to register `donatejohnedwards.com`. Additionally, phishers may simply create subdomains for primary domains that they already own. A phisher who buys the domain `donatefor2008.com`, for example, might simply add DNS records for `johnedwards.donatefor2008.com` and `ronpaul.donatefor2008.com`, among others. These domain names could then be referenced in the phishing emails sent to

From: "Howard Dean" <democraticparty@democrats.org>
Date: August 2, 2007 3:15:53 PM EDT
To:
Subject: They won't steal votes again
Reply-To: dnc-003m8004YD@mailer.democrats.org

THE DEMOCRATIC PARTY

Dear markus,

Last year, we put DNC staff on the ground all across the country to rebuild our party and stand up to the lies and failures of the Republican Party. The 50-State Strategy worked. Those organizers were key to our unprecedented victories up and down the ballot in 2006.

Our organizers are still on the ground in all 50 states, preparing for 2008 in every way possible. And starting this month, they are kicking off an unprecedented voter protection effort, of a scale never attempted by any organization.

While Democrats protect everyone's right to vote, Karl Rove and the Republican Party have a long history of threatening this right, working to make it harder for Americans to vote. We're going to stop them.

Protecting the right of every eligible American to vote is our party's top priority because we know that it's good for America and good for our democracy when everyone votes.

VOTER PROTECTION

help support the organizers working in your state and the vital work they are doing to make sure every vote counts

Contribute

All Americans deserve to go to the polls confident that they won't be harassed or intimidated. That they won't wait hours for a ballot and that their vote will be counted fairly and accurately. Unfortunately, nearly forty-two years after the signing of the voting rights act, millions of American's are treated like criminals just for trying to vote.

Figure 10.7: A portion of a legitimate fundraising email, which allows the recipient to click on the hyperlinked "Contribute" button to support the campaign. This approach would be very easy for a phisher to mimic in an effort to make people submit their credentials to the phisher, thinking they are contributing. Of course, phishers can use inflammatory texts (even more so than political candidates) as calls for action. The authors of this book were able to register the domain `democratic-party.us`, which would be suitable in such an attack, and found a wealth of other cousin name domains available for both parties. Thus, whereas financial institutions typically have registered cousin name domains to defend against abuse, political parties and candidates have not.

potential victims. When clicked on, the link would drive the victim to the fictitious web site.

As we have observed, a significant number of typo domain names have already been registered, or are available to be registered, by parties who are acting in bad faith. Many of these domain names appear so similar to the legitimate domain name that the unsuspecting eye of a potential victim would not notice if directed to one of these sites. Campaigns can take clear and immediate steps to purchase typo domains prior to them falling into the wrong hands. As of this writing, few have done so.

More difficult, however, is the acquisition of cousin domain names. As discussed previously, a significant number of cousin domain names have been registered for both speculative and advertising purposes. Given the near-infinite number of possible cousin domain names, it is unlikely that a campaign could acquire all possibilities. This fact of life provides phishers with the opportunity to register a domain name that may appear similar to the legitimate campaign's web site.

Yet another type of attack might use a spoofed email that appears to come from a political party or candidate to entice recipients to open attachments, thereby infecting their machines with malicious code. Again, this may be done either with the direct goal of spreading malicious code or to deliver a below-the-belt blow to political candidates who rely heavily on the Internet for their communication with constituents.

Even without the registration of a similar domain name, phishers will undoubtedly continue to succeed in constructing emails and web sites that are obvious to detect by a trained eye, but perhaps not so obvious to those who continue to fall victim to them.

10.3 Malicious Code and Security Risks

Malicious code and security risks present some of the more sinister risks to the election process. Malicious code, such as threats that leverage rootkit capabilities,[11] has the potential to gain complete and absolute control over a victim's computer system. Likewise, security risks such as adware and spyware pose serious concerns, both in terms of their invasiveness to a user's privacy (in the case of spyware) and their ability to present users with unexpected or undesired information and advertisements (in the case of adware).

We can consider a number of scenarios where well-known classes of malicious code may be tailored specifically to target those participating in an election. Targets may range from candidates and campaign officials to voters themselves. In discussing these risks we begin with what we consider the less serious category of security risks; we then move into the more serious, insidious category of malicious code.

10.3.1 Adware

Adware, in its truest form, may not pose an immediate and dire risk to the end user. Once installed, however, its control over a user's Internet experience places it into a strategic position on the end user's computer. Adware has the potential to manipulate a user's Internet experience by displaying unexpected or unwanted advertisements. These advertisements may be displayed on the user's desktop or shown to the user through the web browser as the user visits Internet web sites. These advertisements may appear as pop-up windows, or they may appear as content (ads) that are either overlaid or inserted into existing web pages visited by

11. A detailed discussion of rootkits can be found in Chapter 8.

the user. Both techniques have been used frequently by such well-known adware applications as 180Solution's Hotbar [99], Gator Corporation's Gator [96], and WhenU's Save [97]. Adware may be installed by the end user as part of another third-party application, or it may be installed surreptitiously through the exploitation of a software vulnerability in the user's web browser. Chapter 12 discusses adware in more detail.

Adware might be used in numerous ways to influence or manipulate users during the course of an election. In its most innocuous form, adware might simply present the user with advertisements promoting a particular candidate and directing the user to the candidate's web site when clicked. Taking a more deceptive angle, adware might be used to silently replace advertisements for one candidate with another. This may be done directly in the user's browser by manipulating the incoming HTML content before it is rendered or by overlaying a new advertisement on top of an existing advertisement on the user's screen.

Until it is observed in the wild, it is difficult for us to predict the real-world impact that such an adware application might have. It would be important for such an application to be silent and unobtrusive, acting clandestinely to avoid annoying the end user lest its objective backfire. In addition, such an effort may simply help to sway those voters who have not already committed to a particular party or candidate, rather than those voters who have already made their decision.

10.3.2 Spyware

We have frequently seen adware and spyware traits combined into a single application that both delivers advertising and monitors a user's Internet habits. For the purposes of our discussion, we chose to distinguish between the distinct behaviors of adware and spyware, discussing each separately. Spyware, with its ability to secretly profile and monitor user behavior, presents an entirely new opportunity for the widespread collection of election-related trend data and behavioral information.

When discussing the use of spyware, we can conceive of a number of behaviors that might be collected throughout the course of an election in an attempt to provide insight into voters' dispositions. The most basic tactic would be to monitor the browsing behavior of voters and to collect the party affiliations of the Internet sites most frequently visited by the end user. Even without the installation of spyware on an end user's computer, one web site may silently acquire a history of other web sites that the user has previously visited. This capability has been demonstrated by researchers in the past and can be observed at `https://www.indiana.edu/~phishing/browser-recon`. This type of data collection may also include the tracking of online news articles that are viewed and online campaign contributions that are made by determining whether a particular URL was visited.

With the addition of spyware on the end user's computer, these information-gathering efforts can be taken a step further. Emails sent and received by the user can be monitored, for example. In our study, we found that all 19 candidates allow a user to subscribe to their campaign mailing lists, from which a user receives regular frequent updates on the campaign's progress. Knowing how many voters have subscribed to a particular candidate's mailing list may provide insight into the overall support levels for that candidate.

Of course, Internet and browsing behavior alone may not be an indicator of a voter's preference, as voters may be just as likely to visit a competing candidate's web sites and subscribe to a competing candidate's mailing list so as to stay informed about that candidate's campaign. Unfortunately, we could find no prior research that examined the correlation between user Internet behavior and party or candidate affiliation. Nevertheless, spyware clearly poses a new risk in terms of the mass accumulation of election-related statistics that may be used to track election trends.

The collection of voter disposition data is certainly not new, as groups such as the Gallup Organization [137] (known for the Gallup Poll) have been collecting and analyzing user behavior since 1935. What is different in this case is spyware's ability to capture and record user behavior without consent and without the voter's knowledge. Even when a spyware application's behavior is described clearly in an end-user license agreement (EULA), few users either read or understand these complex and lengthy agreements [98]. This changes the landscape dramatically when it comes to election-related data collection.

10.3.3 Malicious Code: Keyloggers and Crimeware

By far one of the most concerning attacks on voters, candidates, and campaign officials alike is that of malicious code infection. Malicious code that is targeted toward a broad spectrum of voters has the potential to cause widespread damage, confusion, and loss of confidence in the election process itself. When we consider the various types of attacks mentioned in this chapter, malicious code—in the form of keyloggers, trojans, and other forms of crimeware—has the potential to carry each of them out with unmatched efficiency. These attacks include the monitoring of user behavior, the theft of user data, the redirection of user browsing, and the delivery of misinformation.

One additional angle for crimeware is the notion of intimidation. Given a threat's presence on a voter's computer, that threat has the potential to collect personal, potentially sensitive information about that individual. This capability may include turning on the computer's microphone and recording private conversations. It may include turning on the computer's video camera and recording activities in the room. It may include retrieving pictures, browser history documents, or copyrighted files from a voter's computer. Perhaps the individual would be turned in to the RIAA if copyrighted music was found on his or her computer. This kind of information gathering creates the potential for an entirely new form of voter intimidation. The collection of such personally sensitive or legally questionable data by a threat might, therefore, allow an attacker to intimidate that individual in an entirely new way. We would, of course, expect and hope that the number of voters who might be intimidated in such a way would be relatively low, but only time will tell whether such speculation becomes reality.

Another form of threat that we have seen in the past involves holding a victim's data hostage until a fee is paid to release it. This possibility was first discussed in [487]. An example of such a threat is Trojan.Gpcoder [340], which encrypts the user's data, erasing the original information, until this fee is paid. Such a threat may present another new form of intimidation whereby the only way for a user to regain access to his or her personal data is to vote accordingly. Such an attack presents obvious logistical challenges. For example, how is the attacker to know which way the victim voted? The attacker may, however, take comfort in the belief that he or she has intimidated enough of the infected population to make a meaningful difference.

Just as the widespread infection of the populace's computers poses a risk to voters, the targeted, calculated infection of specific individuals' computers is equal cause for concern. A carefully placed targeted keylogger has the potential to cause material damage to a candidate during the election process. Such code may also be targeted toward campaign staff, family members, or others who may be deemed material to the candidate's efforts. Such an infection might potentially result in the monitoring of all communications, including email messages and web site access initiated on the infected computer. This monitoring would give the would-be attacker unparalleled insight into the progress, plans, and disposition of the candidate's campaign, perhaps including new messaging, speeches, and otherwise sensitive information critical to the candidate's campaign.

10.4 Denial-of-Service Attacks

Denial-of-service attacks have become increasingly common on the Internet today. These kinds of attacks seek to make a computer network—in most cases, a particular web site—unavailable and therefore unusable. Also known as distributed denial-of-service (DDoS) attacks, they are frequently launched by means of inundating a target with an overwhelming amount of network traffic. This traffic may take the form of Internet protocol requests at the IP and TCP layers or application-level requests that target specific applications such as an organization's web server, email server, or FTP server. Denial-of-service attacks are frequently perpetrated through the use of bot networks, as discussed in more detail in Chapter 7.

A number of high-profile, wide-scale DDoS attacks have demonstrated the effects that such an effort can have. One of the best-known and largest attacks was launched against the country of Estonia in May 2007 [81]. It presented a prime example of a politically motivated attack, as it was launched by Russian patriots in retaliation for the removal of a Soviet monument by the Estonian government. Attackers disabled numerous key government systems during a series of attacks that occurred over the course of several weeks.

In 2006, Joe Lieberman's web site also fell victim to a concentrated denial-of-service attack [397]. Forcing the site offline, the attack paralyzed the `joe2006.com` domain, preventing campaign officials from using their official campaign email accounts and forcing them to revert to their personal accounts for communication.

The implications of such attacks are clear: They prevent voters from reaching campaign web sites, and they prevent campaign officials from communicating with voters.

10.5 Cognitive Election Hacking

Labeled by researchers as *cognitive hacking* [73], the potential for misinformation and subterfuge attacks using Internet-based technologies is as rich as one's imagination. We have already discussed several techniques that may be used to surreptitiously lure users to locations other than a legitimate campaign's web site. These same techniques can be used to spread misleading, inaccurate, and outright false information.

So far, we have discussed typo and cousin domain names that users may visit accidentally when attempting to browse to a legitimate web site. We have also discussed phishing and spam, which have the potential to lure users to web sites by impersonating legitimate candidate web sites. Finally, we have discussed malicious code and the role that it may play in manipulating a user's desktop experience before the user even reaches the intended destination.

The security of a campaign's web site plays another vital role in determining voters' faith in the election process. The breach of a legitimate candidate's web site, for example, would allow an attacker to have direct control over all content viewed by visitors to that web site. This may allow for the posting of misinformation or, worse, the deployment of malicious code to unsecured visitors.

Examples of misinformation about a specific candidate might include a false report about the decision by a candidate to drop out of the race, a fake scandal, and phony legal or health issues. It might also take the form of subtle information that could be portrayed as legitimate, such as a change in a candidate's position on a particular subject, resulting in abandonment of the candidate by voters who feel strongly about that issue.

Attempts to deceive voters through the spread of misinformation are not new. In fact, numerous cases have been documented in past elections using traditional forms of communication [358]. These include campaigns aimed at intimidating minorities and individuals with criminal records, attempts to announce erroneous voting dates, and many other tactics resulting in voter confusion.

During the 2006 election, 14,000 Latino voters in Orange County, California, received misleading letters warning them that it was illegal for immigrants to vote in the election and that doing so would result in their incarceration and deportation. In his testimony before congress, John Trasviña, President and General Counsel of the Mexican American Legal Defense and Educational Fund (MALDEF), discussed this use of misinformation as an example of voter suppression:

> First, the Orange County letter falsely advised prospective voters that immigrants who vote in federal elections are committing a crime that can result in incarceration and possible deportation. This is a false and deceptive statement: Naturalized immigrants who are otherwise eligible to vote are free to vote in federal elections without fear of penalties (including but not limited to incarceration and/or deportation). Second, the letter stated that "the U.S. government is installing a new computerized system to verify names of all newly registered voters who participate in the elections in October and November. Organizations against emigration will be able to request information from this new computerized system." Again, the letter adopts an intimidating tone based upon false information in an apparent attempt to undermine voter confidence within the targeted group of voters. Finally, the letter stated that "[n]ot like in Mexico, here there is no benefit to voting." This letter, representing a coordinated and extensive effort to suppress the Latino vote in the days leading up to a congressional election, has been traced to a candidate running for the congressional seat in the district in which the affected voters live.[12]

12. United States Senate Committee on the Judiciary Prevention of Deceptive Practices and Voter Intimidation in Federal Elections: S. 453 Testimony of John Trasviña. Available at http://judiciary.senate.gov/testimony.cfm?id=2798&wit_id=6514

Another case of deception was targeted at college students in Pittsburgh, Pennsylvania, in 2004 [355]. Canvassers, posing as petitioners for such topics as medical marijuana and auto insurance rates, gathered signatures from students that, unknown to them, resulted in a change to their party affiliation and polling location.

Push polling is one technique that lends itself extremely well to Internet-based technologies. In push polling, an individual or organization attempts to influence or alter the views of voters under the guise of conducting a poll. The poll, in many cases, poses a question by stating inaccurate or false information as part of the question. One well-known push poll occurred in the 2000 Republican Party primary.[13] Voters in South Carolina were asked, “Would you be more likely or less likely to vote for John McCain for president if you knew he had fathered an illegitimate black child?” In this case, the poll’s allegation had no substance, but was heard by thousands of primary voters. McCain and his wife had, in fact, adopted a Bangladeshi girl.

A bill known as the Deceptive Practices and Voter Intimidation Prevention Act of 2007[14] seeks to make these attacks illegal. Currently waiting to be heard in the Senate, it is possible that this bill might be in place for the 2008 federal election, making deceptive tactics such as these illegal, and introducing a maximum penalty of up to 5 years in prison for offenders. This legislation is likely to apply to deceptive practices whether they are performed using traditional communication mechanisms or Internet-based technologies.

While the introduction of such policies is important and provides a well-defined guideline under which to prosecute offenders, only time will tell to what extent legislation will succeed in controlling these acts. As we have seen in some areas, such as the policies developed to outlaw the transmission of spam email, regulations have only marginal effectiveness in reducing the problem. Even today, more than 50% of all email sent on the Internet is purported to consist of spam [401]. There is no reason to doubt that the type of deception and intimidation discussed will be equally successful on the Internet.

The challenge with Internet-based technologies is the ease with which such an attack may be perpetrated. Whereas traditional communication media may have required an organized effort to commit an attack, the Internet allows a single attacker to realize the benefits of automation and scale that previously did not

13. SourceWatch. `http://www.sourcewatch.org/index.php?title=Push_poll`

14. Deceptive Practices and Voter Intimidation Prevention Act of 2007. `http://www.govtrack.us/congress/billtext.xpd?bill=h110-1281`

exist. As such, one person has the potential to cause widespread disruption, with comparably little effort.

Historically, some of the most successful misinformation attacks on the Internet have been motivated by profit. Pump-and-dump schemes [369], for example, have become an extremely common form of spam. These schemes involve the promotion of a company's stock through the issuance of false and misleading statements. After the stock price rises owing to renewed interest from the message's recipients, the perpetrators sell their own stock for a substantial profit.

One significant surge of pump-and-dump emails that was observed in 2006 was attributed to a bot network, operated by Russian fraudsters [268]. In this attack, 70,000 infected computers spread across 166 countries were organized into a bot network that was used to send out unsolicited stock-promoting spam. Such a network could easily be directed to send any form of email, including disinformation and fallacies related to a candidate, voters, and the election itself. Chapter 7 discusses botnets and their applications in more detail.

10.6 Public Voter Information Sources: FEC Databases

The Federal Election Commission [62] was created both to track campaign contributions and to enforce federal regulations that surround them.

> In 1975, Congress created the Federal Election Commission (FEC) to administer and enforce the Federal Election Campaign Act (FECA)—the statute that governs the financing of federal elections. The duties of the FEC, which is an independent regulatory agency, are to disclose campaign finance information, to enforce the provisions of the law such as the limits and prohibitions on contributions, and to oversee the public funding of Presidential elections.

To provide a public record of campaign contributions, the FEC must maintain, and provide to the public, a full record of all campaign contributions. Many web sites that allow online contributions clearly indicate their requirement to report those contributions to the Federal Election Commission. The following text, taken from one candidate's web site exemplifies this kind of disclaimer:

> We are required by federal law to collect and report to the Federal Election Commission the name, mailing address, occupation, and employer of individuals whose contributions exceed $200 in an election cycle. These records are available to the public. However, they cannot be used by other organizations for fundraising. We also make a note of your telephone number and email address, which helps us to contact you quickly if follow-up on your contribution is necessary under Federal election law. For additional information, visit the FEC website at `http://www.fec.gov`.

The FEC's role is to make this data available to the public. The information is available as raw data files, via FTP, and through online web interfaces on the FEC web site.

Numerous third-party web sites, such as `http://www.opensecrets.org`, also use this data to provide regular high-level reports on candidate funding. Consumers of the data are restricted by a policy that regulates how the data can be used [64]. The policy is surprisingly lenient, as it is primarily intended to prevent the use of contributors' names for commercial purposes or further solicitation of contributions.

The information provided in this database consists of each contributor's full name, city, ZIP code, and particulars of the contribution, such as the receiving candidate or party, the amount, and the date of the contribution. While limited, this information does allow one to build a history of political contributions for any U.S. citizen who appears in the database.

Contributors of record may be more likely to become victims of the other attacks already discussed in this chapter. Appearing in this database may expose high-net-worth contributors to targeted phishing (spear phishing) or malicious code attacks if the individual's name can be connected to his or her email address (no longer a difficult feat).

10.7 Intercepting Voice Communications

While this chapter has focused primarily on Internet-based risks, we would be remiss if we did not discuss at least one additional risk given a recent particularly noteworthy and sophisticated attack against a foreign nation's communication infrastructure. Labeled the *Athens Affair* by authors Vassilis Prevelakis and Diomidis Spinellis [320], this well-coordinated attack highlighted the increased role that common technologies play in all forms of our daily communications. In their paper, the authors retrace the alarming events related to the interception of cell phone communications from high-ranking Greek government officials:

> On 9 March 2005, a 38-year-old Greek electrical engineer named Costas Tsalikidis was found hanged in his Athens loft apartment, an apparent suicide. It would prove to be merely the first public news of a scandal that would roil Greece for months.
>
> The next day, the prime minister of Greece was told that his cell phone was being bugged, as were those of the mayor of Athens and at least 100 other high-ranking dignitaries, including an employee of the U.S. embassy.
>
> The victims were customers of Athens-based Vodafone-Panafon, generally known as Vodafone Greece, the country's largest cellular service provider; Tsalikidis was in charge of network planning at the company. A connection seemed obvious. Given the

> list of people and their positions at the time of the tapping, we can only imagine the sensitive political and diplomatic discussions, high-stakes business deals, or even marital indiscretions that may have been routinely overheard and, quite possibly, recorded.
>
> Even before Tsalikidis's death, investigators had found rogue software installed on the Vodafone Greece phone network by parties unknown. Some extraordinarily knowledgeable people either penetrated the network from outside or subverted it from within, aided by an agent or mole. In either case, the software at the heart of the phone system, investigators later discovered, was reprogrammed with a finesse and sophistication rarely seen before or since.

In this attack, perpetrators used rootkit techniques, like those discussed in Chapter 8, on the cellular provider's phone switch to remain hidden. Over the past two decades, the basic communications systems that we rely on for both our traditional land-line telephones and our cellular phone communications have increasingly moved to commodity-based hardware and software [108]. In the past, would-be attackers were forced to learn complex and proprietary embedded systems, making the introduction of malicious code on these systems difficult, if not impossible. Today's commoditization simplifies this effort, as witnessed by the attack discussed here, and greatly increases the likelihood that an attacker might gain a similar foothold on communications systems in the future.

Central switching networks are not the only target. Mobile devices themselves remain even more likely candidates for interception of communications. Today's mobile devices, an increasing number of which can be considered smartphones, provide ripe opportunities for the introduction of malicious code. While traditional threats such as viruses, worms, and trojans have yet to gain widespread prominence on mobile devices (although they do exist), the potential for targeted customized mobile threats has existed for some time.

One particular application, known as FlexiSpy and sold by Bangkok, Thailand–based software vendor Vervata, allows listening to a remote phone's surrounding while it is not in use (Figure 10.8). It also allows retrieval of the phone's personal data and monitoring of all email and SMS messages sent by the phone. The software itself is available in "Pro," "Light," "Alert," and "Bug" versions. The vendor prides itself on its software's ability to remain hidden and unnoticeable on an infected device.

The infection of a candidate, campaign staff, or candidate's family's cell phone with such a freely available application could have dire consequences. All back-room and hallway conversations engaged in by the candidate could be monitored at all times and intercepted by the attacker. Worse, opinions—perhaps including those not shared with the public or outsiders—could be recorded and

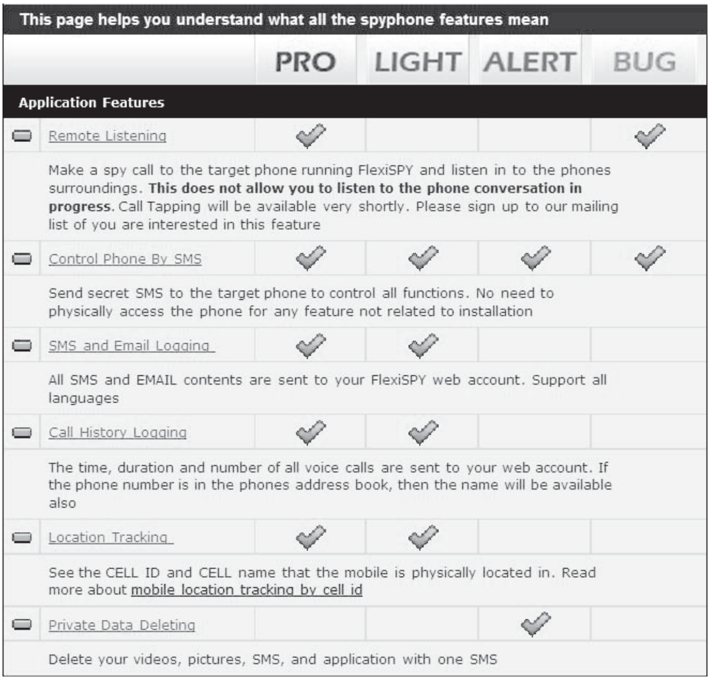

This page helps you understand what all the spyphone features mean

Application Features	PRO	LIGHT	ALERT	BUG
Remote Listening	✓			✓
Make a spy call to the target phone running FlexiSPY and listen in to the phones surroundings. **This does not allow you to listen to the phone conversation in progress.** Call Tapping will be available very shortly. Please sign up to our mailing list of you are interested in this feature				
Control Phone By SMS	✓	✓	✓	✓
Send secret SMS to the target phone to control all functions. No need to physically access the phone for any feature not related to installation				
SMS and Email Logging	✓	✓		
All SMS and EMAIL contents are sent to your FlexiSPY web account. Support all languages				
Call History Logging	✓	✓		
The time, duration and number of all voice calls are sent to your web account. If the phone number is in the phones address book, then the name will be available also				
Location Tracking	✓	✓		
See the CELL ID and CELL name that the mobile is physically located in. Read more about mobile location tracking by cell id				
Private Data Deleting			✓	
Delete your videos, pictures, SMS, and application with one SMS				

Figure 10.8: FlexiSpy, developed and sold by Bangkok, Thailand's Vervata, allows for monitoring and tapping of cell phone communications. It is supported on Windows Mobile, Symbian OS, and Blackberry devices. Today installation requires physical access to the device. Much like desktop operating systems, however, future versions might be installed through software vulnerabilities or messaging applications.

made available for later playback, introducing the potential for widespread exposure and damage.

We have already seen examples of unexpected recordings accidentally made public for other political figures, including those involving California Governor Arnold Schwarzenegger in 2006 and 2007 [308]. In that case, the recordings were unintentionally exposed through the governor's web site and resulted in criticism of a number of his comments that were made without the intent of them becoming public.

Conclusion

As campaigns increasingly look to the online medium to gather support, it is important to consider the inherent risks that will follow. In this chapter, we discussed a number of risks that may present themselves in any election campaign; however, it is important to acknowledge that many more remain that we have not discussed.

It is apparent both from past events and from our findings that candidates and their campaigns are just beginning to understand the risks of online advocacy and have yet to take the necessary precautions to protect themselves. Our fear is that a true appreciation of the required countermeasures will not be realized until these attacks do, in fact, manifest themselves.

Many of these individual risks, when combined, would result in increasingly sophisticated attacks. While we have discussed many of these risks independently, the combination of these threats creates complex new variations that are already being seen in the wild in other areas such as online banking and ecommerce.

Our goal in writing this chapter was not to sow seeds in the minds of would-be attackers or to spread fear, uncertainty, and doubt, but rather to discuss real-world risks that already exist. None of the attacks discussed here are new or novel; we have simply applied them to a specific recurring event, the election process. Our hope is to raise awareness of the potential risks before they are able to manifest themselves in the upcoming 2008 federal election, or any election that follows.

One thing is clear: It is impossible for us to predict how successful any one of these attacks might be in making a material impact on the election process. Given our experiences with previous widespread Internet-borne risks, we certainly do have an appreciation and respect for the potential that they present. While that is not to be discounted lightly, only time will tell how dangerous they become.

In addition, if a successful widespread attack were to occur (one that was recognized to have swayed the vote), what recourse is there? What if intimidation, misinformation, and infectious election-targeted malicious code become the norm?

Acknowledgments

The author wishes to thank Markus Jakobsson for providing a political fundraising email sample, for pointing out additional political-sounding URLs that he had registered, and for contributing helpful discussions, feedback, and comments on earlier drafts. In addition, the author would like to thank Zulfikar Ramzan for his advice, feedback, and recommendations during the writing of this chapter, and Kevin Haley for his recommendation on the addition of push polling.

Index

J

K

N

O

P

T