Clear, unambiguous password policies: An oxymoron?

Michelle Steves[1], Kevin Killourhy[2], Mary Theofanos[1]

[1]National Institute of Standards and Technology, Gaithersburg, Maryland, United States
{michelle.steves, mary.theofanos}@nist.gov
[2] formerly of the National Institute of Standards and Technology, Gaithersburg, Maryland, United States

Abstract. Password policies – documents which regulate how users must create and manage their passwords – can have complex and unforeseen consequences on organizational security. Since these policies attempt to govern user behavior, users must be clear as to what is expected of them for a policy to be effective. While a culprit of misinterpretation, policy ambiguity also prevents researchers from comparing and contrasting policy statements. To tackle ambiguity, we developed a formal language for stating what behavior is and is not allowed when creating, managing, and changing passwords. This formal language lends itself to policy analysis and visualization. A corpus of 41 password policies was translated into the formal language and analyzed. Having these clear, unambiguous policy statements enables us to explore password policies in much greater detail, discuss the relative merits of different statements, compare and contrast policies, and begin to examine the interplay between usability and security in password policies.

Keywords: Usable security, password policy, formal language, Extended Backus-Naur Form

1 Introduction

As part of the Research and Development thrust of the United States' Comprehensive National Cyber-Security Initiative, we undertook an exploration of the relationship between usability and security in password policies. In this domain, as in many areas of computer security, we find a mixture of security concerns and usability concerns. Adams and Sasse [1] observed that when users feel that password rules are preventing them from accomplishing their work, they will work around them. For instance, if files cannot be easily shared among account owners, users will share an account, violating the password policy restriction on sharing. Usability issues related to security policies affect security-related behaviors and by extension, security outcomes.

Password policies are documents that specify how passwords must be created, stored, changed, and managed. Depending on the security concerns of the policy writer, these documents regulate everything from length and lifetime, to when and how a user can transmit passwords. For example, the SANS Institute [16] has shared a template password policy for use by policy writers. The template requires

adfa, p. 1, 2011.

that user-account passwords be changed every six months, and it offers guidance on what constitutes a strong password (e.g., "at least fifteen alphanumeric characters"). The template requires that users not reveal their password, not check "remember password" boxes in applications, and not store passwords in written form.

These policy statements are created with security concerns in mind. The restriction on users storing their password in writing is meant to protect the password, but it requires user "buy in" if that protection is to exist. Consequently, the role of the user cannot be discounted when developing a policy. Since some security requirements are all but unenforceable, security depends on a policy that can effectively explain what is expected of a user. Unfortunately, security concerns that prompt statements in a password policy are rarely understood by users. Typical policy requirements include a password with an 8 character minimum length (or 15 in the example above), comprised of mixed-case letters, numbers, and special characters, and which does not contain dictionary words or personal information. Cheswick [3] has called these "eye of newt" policy statements in part because they sound like a magic formula to a user.

A typical user is likely governed by multiple policies both at work and at home (e.g., to access corporate email servers, personal financial information, and e-commerce sites). Ambiguities in these policies, discrepancies among them, and the sheer number of different policies may cause confusion. As a result of this cognitive burden, users may choose weak passwords, write them down, or violate policies in other ways. As policy violation becomes an accepted matter of course, the security goals of the policy are not met. Consequently, overall security may be weakened.

2 Problem and Approach

In a preliminary survey of password policies, we discovered an unexpected problem both for our research effort and for users: ambiguity. We were surprised at the frequency with which there was disagreement over what a policy statement meant. If we had difficulty agreeing, would users who are expected to follow these policies also have different interpretations? An example of the difficulty can be seen by examining a statement in the password policy of one Federal agency:

Passwords contain a combination of letters, numbers, and at least one special character.

Different people have interpreted this statement to mean each of the following:

1. Users must create passwords with at least one letter, at least one number, and at least one special character (i.e., defining a combination to mean one or more of each).
2. Users must create passwords with at least one special character (i.e., defining a combination to mean zero or more of each, except special characters since at least one is explicitly required).
3. Users must create passwords with at least two letters, two numbers, and one special character (i.e., because letters and numbers are both plural).

All interpretations are reasonable and as a result a reader cannot be sure whether his or her interpretation is correct.

Ambiguity is both a usability issue and a security issue. Users faced with such a policy cannot know which passwords are in compliance and which are not. Users who take the time to create a memorable, strong password only to have it rejected as non-compliant, are unlikely to try as hard the second time. Researchers cannot accurately assess security when there is disagreement over the security properties. For instance, each of the three interpretations corresponds to a password space of different size; the strength of the password depends on the size of this space.

3 Background

Our effort exists within 30 years of research into password usability and password policies. Research studies have periodically tried to characterize the passwords that people choose [4], [6], [11], [14], and [20] to name a few. These works usually rely on a list of passwords obtained by a system-administrator, observed through a toolbar, or revealed through a security compromise. They provide empirically grounded insight into the actual passwords that people choose. While these studies do not consider password policies directly, they are related to our research on policy since users' password choices are governed in part by policy requirements.

Some researchers have attempted to examine the connection between password policy and user behavior already. Mannan and Oorschot [13] surveyed users of online banks regarding their understanding of bank's security requirements. They found a disconnect between bank guidelines and user practice. For instance, bank password-change recommendations were on the order of a few months, yet the majority of users do not change their passwords within a year. Furnell [8] surveyed 10 popular websites, examining the password-creation guidelines, the enforcement of password-composition restrictions, and the reset policy. Inglesant and Sasse [9] surveyed 32 staff members at a research university and a financial-services organization about their password usage. They found that password policies which ignore human factors may result in unexpectedly poor security (e.g., shared and written-down passwords).

Many authors have observed the increasing burden of password management on users. Summers and Bosworth [19] argue for user guidance and enforcement to prevent users from choosing easily guessable passwords. Spafford [18] argued that the best practices shared by many modern password policies are actually artifacts based on out-of-date risk assessments. For instance, he traces requirements that passwords be changed monthly to the computing power of Department of Defense machines in the 1970s. He advocates updating these out-of-date best practices to modern, sound practices based on individual risk assessments. Farrell [5] introspectively determined that he used over 61 passwords, 15 of which were committed to memory. He argues that policy writers must acknowledge the increasing burden of password management on their users.

Some research has attempted to make writing security policies easier. A visualization of access-control policies is presented in [21]. An extensive list of guidelines for

anyone architecting a system for writing security and privacy policies is provided in [10].

Some research has focused specifically on writing password policies. A language for expressing a password-policy scenario in a formal language is presented in [17]. Using simulation, a measure of system harm resulting from the given scenario can be estimated. An ontological framework for reasoning about the security and usability costs of different policy decisions is presented in [15]. These works show the potential for password policies written in formal language. That effort partners well with the aim of our current work to bridge the gap between ambiguous, informal policy statements and clear, explicit, formal policy statements.

Perhaps the closest research efforts were those documented in [2] and [7]. Bonneau and Preibusch [2] surveyed the empirical password policies of 150 websites. By creating accounts and systematically changing the password, they were able to infer the enforced lengths and complexity requirements. Florêncio and Herley [7] surveyed length and complexity requirements for 75 websites according to their password policies. Their primary finding was that length and complexity requirements are not explained by the size of the assets protected but by the dependence on advertising dollars. In other words, sites that depend on users for revenue have weaker, more accommodating policies.

While our research overlaps that of [2] and [7], the focus of our work differs in important ways. Foremost, our interest in capturing the diversity and ambiguities in password policies reflects a different research goal. Our investigation moves beyond length and complexity requirements and into the numerous other requirements and guidance that password policies provide.

4 Creating a Taxonomy of Policy Statements

We created the taxonomy by first analyzing a corpus of collected password policies. Then, we developed the grammar whereby the regulations of the password policies could be expressed as explicit expectations on user behavior when creating, changing, storing, or communicating passwords. Next, we used this grammar to translate the 41 password policies into a formal, well-defined language. Finally, we explore how policies, once represented in this language, lend themselves to comparison, visualization, and further analysis. Each of these steps in our method is described in the following sections.

4.1 Collecting a Policy Corpus

We employed two methods for collecting password policies for study. We collected workplace policies by searching Federal Government websites. We collected personal-site policies by searching popular websites. The two approaches enabled us to collect 41 policies: 22 concerning workplace passwords, and 19 concerning personal-site passwords.

In an effort to collect workplace policies, we visited each of the Federal Departments and Agencies listed at www.usa.gov (August 2010). On each site, we attempted to find the password policy. To qualify, the policy document itself had to be obtained. Policies were also checked to ensure that they applied to employees of the organization (rather than outside users of a service provided by the organization). Policies judged to govern workplace behavior concerning passwords were added to the corpus. From the sites visited, we obtained 22 policies.

To collect a comparatively sized set of policies from sites that users might visit for personal use, we began with the corpus of online password policies assembled in [7]. This paper explains the collection methodology in detail, and importantly, one aim was a representative collection. In [7], 77 policies were examined and links to 56 of them were provided. We visited all 56 links. If a link was broken or we could not find a minimal set of policy rules, we excluded the policy from our corpus by necessity. For a policy to be included in our corpus it must have, at a minimum, a requirement or recommendation about password length or a statement about which character sets are acceptable. Following this procedure, we added 15 policies to our corpus. Because our corpus contained a paucity of financial-service policies, and because we felt that such policies would be important in our corpus (as they are policies from personal-use sites that are likely to have detailed security concerns), we visited the top business and financial services websites, ranked by US traffic, according to www.alexa.com (in August 2010). We employed the same strategy to find a password policy governing users of these sites as above. Following this procedure, we added four policies, for a total of 19 personal-use policies and 41 policies overall.

4.2 Developing a Taxonomic Grammar

While reading and becoming familiar with the regulations and guidelines of the policies in the corpus, we made several attempts at organizing the policies and policy statements into rules. Our final product was a formal language in which policies were represented as a set of statements formed from a subset of English. However, in the process of arriving at this approach, we attempted two other approaches: feature vectors (an n-dimensional vector of features to represent a policy) and qualitative coding.

Our first attempt at organizing password policies was to encode them as feature vectors. Earlier research, such as that by Florêncio and Herley [7], effectively encoded password length, composition requirements, and lifetime in this manner. However, because our interest was in capturing other aspects of policy behavior (e.g., what word and letter combinations are forbidden, and whether users can use password vaults), this approach quickly became unwieldy.

Our second attempt at organizing policy statements was to analyze password policies in the style of a social scientist trying to organize research data. Using a software package for qualitative data analysis, a researcher tagged sentences and sections in a password policy with a topic code. For instance, a line insisting that strong passwords have a minimum of 15 characters would be coded “Length,” and a line stating that passwords must be alphanumeric would be coded “Character Content.” We identified 9 top-level codes, in this way: Length, Lifetime, Lockout, Storage, Use, Character

Set, Character Content, Content Features, and Forgotten. The plan was to subdivide each of these into different kinds of statements within each topic area. This approach also became unwieldy, as the distinction between topic areas was subject for debate.

These failed attempts contributed to our final, more successful, approach to studying password policies. Rather than trying to divide the password policy into topic areas, as in the previous approach, we simply translated each policy statement into clear language that explicitly states what is expected of the user. For instance, "strong passwords are at least 15 characters and alphanumeric" became two statements: (1) "users must create passwords that are at least 15 characters long" and (2) "users must create passwords where every character is a letter or a number." Note that the user is the subject of the sentence, making it clear that the policy applies to the user (compare "passwords must be stored as cryptographic hashes" which is probably, but not explicitly, regulating the behavior of a system administrator, not a user).

After rewriting several policies, we discovered that some sentences were being reused. Reviewing what we had written, we observed that other sentences were saying the same thing but in different ways. By restricting ourselves to a grammar in which a policy statement could only be expressed in one way – and two different statements imposed two different requirements, even if only subtly – the idea of a formal grammar for expressing password policies arose.

The actual grammar of the formal language was specified using an extension of the Backus-Naur Form (EBNF). The language and the procedure for translating a policy into the language is described in [11], as well as, both the EBNF syntax and a description of the semantics.

4.3 Translating Policies to the Formal Language

Translation involved two steps. First, sentences, bullets, or policy statements that were within scope were identified. Then, each statement was translated to one or more statements in the formal language. Each of these steps is described in much greater detail in [11], but the procedure is fairly straightforward. The translated policy documents were checked with a language parser to ensure that the syntax of each statement was correct. In the end, all 41 policies in the corpus were translated into the formal language.

4.4 Analysis and Visualization

The translated policy documents enable analysis and visualization that would not be possible with the raw documents. Specifically, since the formal grammar has the property that semantically equivalent statements are syntactically equal, we can analyze statement prevalence and frequency. Further, since the language has a hierarchical characteristic whereby similar statements start the same way, a tree-like visualization of password policies becomes possible.

Policy Summary Statistics. Some summary statistics provide a clear picture of the diversity of password policies. The 41 policies yielded 449 statements in total. The

policy with the fewest statements had only one: *Users must create passwords with length greater than or equal to 6 characters.* The policy with the most statements had 38. The median number of statements per policy was 9.

Of the 449 total statements, there were 153 unique statements, meaning that there is overlap across policies. The statement which appeared in the most policies was *Users must create passwords with length greater than or equal to 8 characters.* This statement appeared in 23 policies (56%). The next most prevalent statement appears in only 15 policies (37%): *Users must not communicate passwords to anyone*. Regarding the infrequent statements, 71 of the 153 statements (46%) appeared in only one policy. Additionally, of the 41 policies, no two policies shared the exact same set of statements.

Visualization with Policy Trees. One powerful advantage of the taxonomy of password statements is a visualization based on the parse tree. Figure 1 shows an example of such a visualization. Each statement can be broken up into phrases based on which grammar rule matched the words. For instance, the policy statement "*Users* **|** *must* **||** *create passwords* | *with length* | *greater than or equal to* | *8 characters*" has vertical bars delimiting the phrases. The double vertical bar indicates a blank phrase (contrasting with the phrase "*not*" which can appear instead of the blank).

We can visualize each policy statement as a series of line segments, one per phrase, which radiate out from a central point. If two policies start with the same phrases, they follow the same path. When the two policies use different phrases, the paths diverge. Since a policy is a collection of statements, we can view it as a tree.

Figure 1 shows a tree with 153 paths from the center of the circle to the outside, each corresponding to a different one of the 153 unique statements in the corpus. The darker paths correspond to policy statements made by the (now obsolete) National Institute of Standards and Technology (NIST) password policy. The lighter paths correspond to statements that are in other policies and not NIST's. The arcs around the outside of the figure reflect that different slices (or wedges) of the circle correspond to different prefix phrases. The median number of phrases is 6 in a policy, with a minimum of 4 ("Users | must | not | communicate passwords | except in an emergency") and a maximum of 7 ("Users | must || create passwords || in the set of | strings with | at least 2 unique characters"). Note that, by convention, since all policies start with the phrase "Users," we do not count it when counting the number of phrases in a statement.

We believe the visualizations provide a useful indication of policy complexity and structure. Policies with more statements produce denser trees. We also find these visualizations to be useful in comparing policies by placing trees for two policies side-by-side. Since policy statements correspond to the same path relative to the side-by-side centers, it is possible to spot paths that are shared and paths that are not. It is also possible to compare regions of coverage (e.g., which policies have many strict "must" regulations and which policies have more suggestive "should" guidelines).

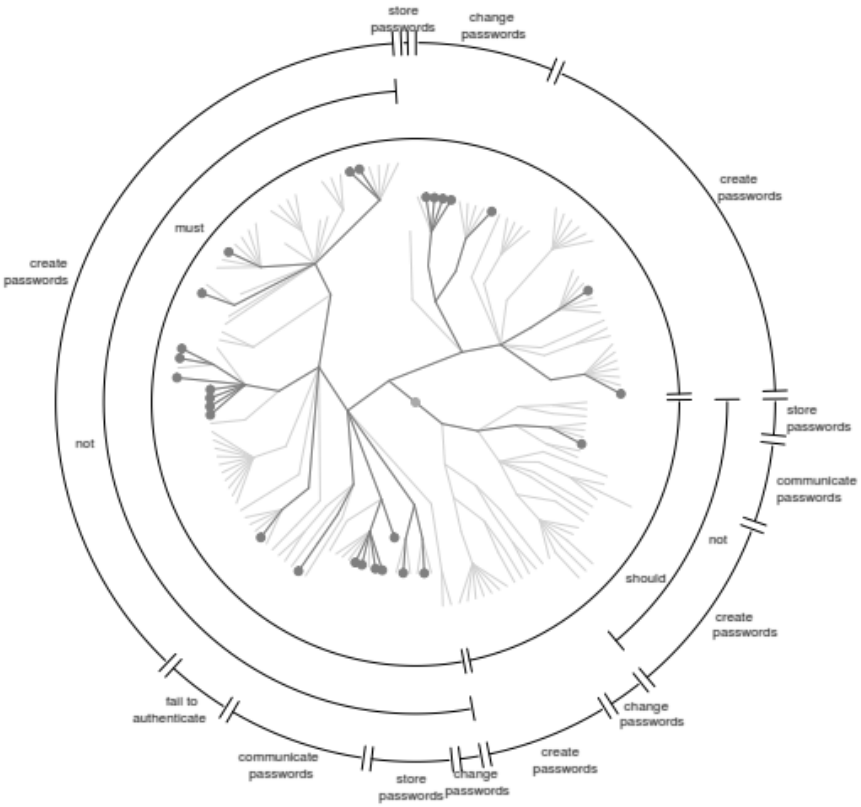

Fig. 1. Password policy tree-like visualization

5 Explorations

To highlight the potential of this taxonomy of password-policy statements, we conducted several exploratory investigations. First, we compared policies that have a known relationship to see how that relationship is reflected in the visualization. Second, we compared the workplace policies to the personal-site policies to see whether workplace policies are stricter (as would be expected based on earlier results by Florêncio and Herley [7]).

5.1 Comparing Policies: NIST, the Census Bureau, DOC and the Marines

The password-policy visualizations make it easy to explore and compare different password policies. Figure 2 shows four policies in a 2 × 2 tiling. The policies on the top row are from NIST and the Census Bureau. The policies on the bottom line are from the Department of Commerce (DoC) and the Marine Corps Enterprise Network

(MCEN). Note that NIST and the Census Bureau are both organizations within the Department of Commerce.

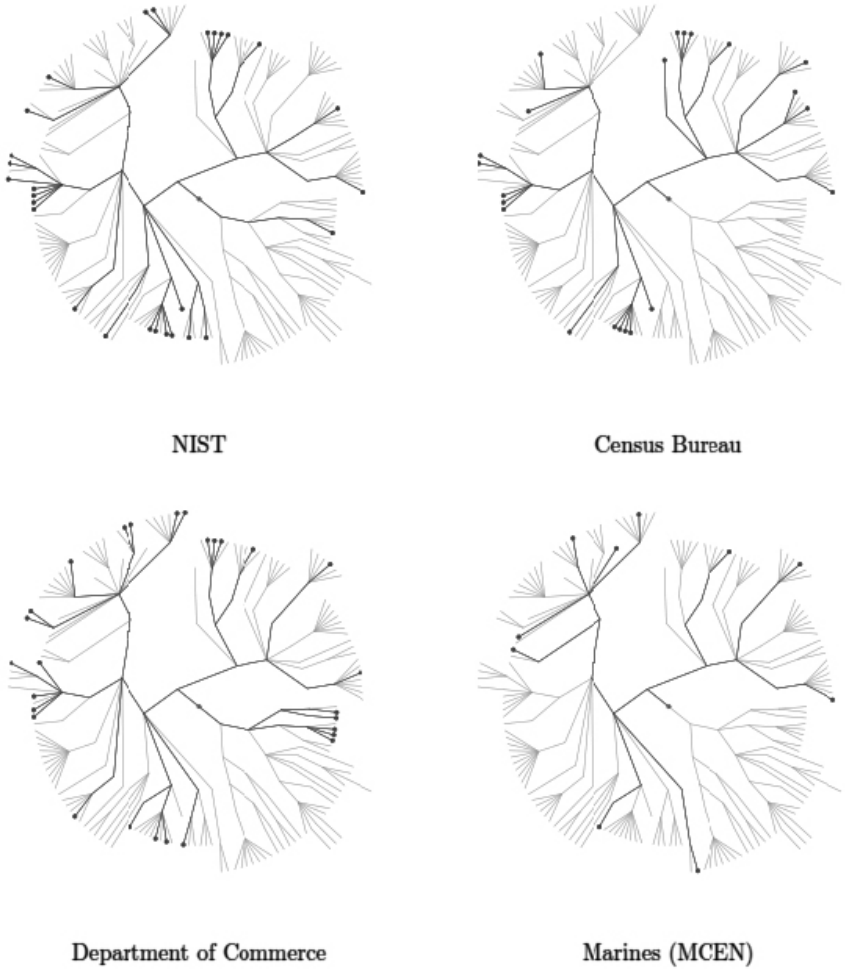

Fig. 2. Panels comparing four different workplace password policies

Given their department-level connection, one might expect to see similarities between the NIST and Census policies. Indeed, we observe similarities, and also some differences. For instance, between 12 o'clock and 1 o'clock are a set of restrictions on when passwords must be immediately changed. Specifically, the prefix is "Users must change passwords immediately if" and there are five different completions, corresponding to the path's split into a five-tined fork. Neither NIST nor Census includes the first completion: sent unencrypted (i.e., a requirement that passwords be changed immediately if sent in clear text). Both NIST and Census include the next three completions: found non-compliant, directed by management, and compromised. NIST

also includes the final completion while Census does not: shared. NIST's policy allows passwords to be communicated to others in an emergency, but this statement tries to ensure that shared passwords are changed as soon as possible.

The similarities between the two policies suggest a common source. The two policies share many of the same statements with the password policy for the Department of Commerce, shown in the lower left of Figure 2. The DoC policy shares the three regulations on immediate password change: found non-compliant, directed by management, and compromised. However, there are some discrepancies between the department-level policy and those of the organizations in the department. For instance, in the 3 o'clock position, both NIST and Census have the same path while DoC has a slightly different one (i.e., NIST and Census have a path that angles downward in the last phrase while DoC has a path that angles slightly upward. This corresponds to an 8 character minimum for NIST and Census and a 12 character minimum for DoC.

In contrast to the NIST, Census, and DoC policies which have dozens of statements, the Marine (MCEN) password policy in the lower right is less densely populated. It is comprised of only 10 statements. This policy says nothing about the conditions under which passwords must be immediately changed. Also, in the 6 o'clock position, the MCEN policy has a rule that is in none of the other three: Users must not change passwords before 7 days. The MCEN policy is one of a few that set a limit on how often a user can change his or her password (presumably to prevent them from returning too quickly to an old password). This exploration of the common and the unusual across policies promotes discussion of what statements should be in a policy.

We also found it interesting that, despite their differences, the NIST, Census, and DoC policies look similar, especially when compared with a policy from unrelated organizations in the Federal government.

5.2 Comparing Policy Groups: Workplace and Personal-Site

In addition to comparisons of individual policies, we can compare sets of policies by comparing the statements which appear in each set. For instance, one might want to compare workplace and personal-site policies to see whether workplace policies tend to contain more statements containing requirements. Florêncio and Herley [7] found that sites which draw ad-revenue have more lenient policies, and sites which do not need to entice their users to join (i.e., workplaces) are stricter.

A statistical analysis of the number of statements in each policy reveals that the median for workplace policies was 10 statements, and the median for personal-site policies was 7. Moving beyond summary statistics, it is possible to visually contrast the workplace policy tree and the personal-site policy tree. The workplace tree is comprised of the union of all the workplace statements, and the personal-site tree is comprised of the personal-site statements (not shown due to space limitations).

Partitioning of policies into workplace and personal-site is only one possible partitioning of many. Since the personal-site policies are comprised of sites ranging from banks to sports news, we expect that a subdivision of the personal-site policies would be informative. One might imagine that banks and other financial services companies would have more restrictive policies as security is more of a concern. The formal-

language representation and visualization lends itself to such partitioning, exploration, and analysis.

6 Summary and Future Work

This work arose from a need to deal with ambiguities in current password policies. These ambiguities not only prevent researchers from comparing and contrasting policy statements, they can also cause users to misinterpret what is expected of them. To tackle ambiguity, we developed a formal language for stating what behavior is and is not allowed when creating, managing, and changing passwords. This formal language lends itself to policy analysis and visualization.

We have used this language to begin to explore the diversity and complexity of password policies by examining a corpus of password policies. Having clear, unambiguous policy statements enables us to study password policy statement properties. What quantifiable security benefits are provided by a statement or set of statements? What usability issues are introduced? These usability issues can also be quantified and studied empirically in order to understand the security and usability interplay.

In addition to the study of policy statement properties, the formal language enables us to examine how users interpret individual policy statements and develop a mapping of the formal policy statements to plain language statements that remains unambiguous, minimizes misinterpretation, and enhances understanding by users.

References

1. Anne Adams and Martina Angela Sasse. Users are not the enemy. *Communcations of the ACM*, 42(12):41–46, 1999.
2. Joseph Bonneau and Sören Preibusch. The password thicket: Technical and market failures in human authentication on the web. *9th Workshop on the Economics of Information Security (WEIS-2010)*, June 7–8, 2010, Cambridge, MA, 2010. `http://weis2010.econinfosec.org/papers/session3/weis2010_bonneau.pdf`, accessed Jan. 2014.
3. Bill Cheswick. Rethinking passwords. Presentation at the Solaris Security Summit, Baltimore, MD, 2010. `http://www.cheswick.com/ches/talks/baltimore.pdf`, accessed Jan 2014.
4. Matteo Dell'Amico, Pietro Michiardi, and Yves Roudier. Password strength: An empirical analysis. In *Proceedings of 29th IEEE Conference on Computer Communication (INFOCOM-2010)*, pages 1–9, Mar 14–19, 2010, San Diego, CA, 2010. IEEE Press.
5. Stephen Farrell. Password policy purgatory. *IEEE Internet Computing*, 12(5):84–87, 2008.
6. Dinei Florêncio and Cormac Herley. A large-scale study of web password habits. In *Proceedings of the 16th International World Wide Web Conference (WWW-2007)*, pages 657–666, May 8–12, 2007, Banff, Alberta, 2007.

7. Dinei Florêncio and Cormac Herley. Where do security policies come from? In *Proceedings of the 6th Symposium on Usable Privacy and Security (SOUPS-2010)*, July 14–16, 2010, Redmond, WA, 2010. ACM Press, New York, NY.
8. Steven Furnell. An assessment of website password practices. *Computers & Security*, 26:445–451, 2007.
9. Philip G. Inglesant and M. Angela Sasse. The true cost of unusable password policies: Password use in the wild. In *Proceedings of the 28th International Conference on Human Factors in Computer Systems*, pages 383–392, Apr 10–15, 2010, Atlanta, GA, 2010. ACM Press, New York, NY.
10. Maritza Johnson, John Karat, Klare-Marie Karat, and Keith Grueneberg. Optimizing a policy authoring framework for security and privacy policies. In *Proceedings of the 6th Symposium on Usable Privacy and Security (SOUPS-2010)*, July 14–16, 2010, Redmond, WA, 2010. ACM Press, New York, NY.
11. Kevin Killourhy, Yee-Yin Choong, and Mary Theofanos. Taxonomic rules for password policies: translating the informal to the formal language. National Institute of Standards and Technology, Gaithersburg, Maryland, NISTIR 7970, Dec 2013.
12. Daniel V. Klein. Foiling the cracker; a survey of, and improvements to unix password security. In *Proceedings of the 2nd USENIX Security Symposium*, pages 5–14, Aug 27–28, Portland, OR, 1990. USENIX.
13. Mohammad Mannan and Paul C. Van Oorschot. Security and usability: The gap in real-world online banking. In *Proceedings of the New Security Paradigms Workshop (NSPW-2007)*, pages 1–14, Sep 18–21, 2007, North Conway, NH, 2007. ACM Press, New York, NY.
14. Robert Morris and Ken Thompson. Password security: A case history. *Communications of the ACM*, 22(11):594–597, 1979.
15. Simon E. Parkin, Aad van Moorsel, and Robert Coles. An information security ontology incorporating human-behavioural implications. In *Proceedings of the 2nd International Conference on Security of Information and Networks (SIN-2009)*, pages 46–55, Oct 6–10, 2009, Famagusta, North Cyprus, 2009. ACM Press, New York, NY.
16. SANS Institute. SANS password policy, 2006. `http://www.sans.org/security-resources/policies/Password_Policy.pdf`, accessed Jan. 2014.
17. Richard J. K. Shay, Abhilasha Bhargav-Spantzel, and Elisa Bertino. Password policy simulation and analysis. In *Proceedings of the ACM Workshop on Digital Identity Management*, pages 1–10, Nov 2, 2007, Fairfax, VA, 2007. ACM Press, New York, NY.
18. Gene Spafford. Security myths and passwords, 2006. `http://www.cerias.purdue.edu/site/blog/post/password-change-myths/`, accessed Jan. 2014.
19. Wayne C. Summers and Edward Bosworth. Password policy: The good, the bad, and the ugly. In *Proceedings of the Winter International Symposium on Information and Communication Technologies*, pages 1–6, Jan 5–8, 2004, Cancun, Mexico, 2004. Trinity College, Dublin.
20. Thomas Wu. A real-world analysis of Kerberos password security. In *Proceedings of the ISOC Symposium on Network and Distributed System Security (NDSS-1999)*, San Diego, CA, 1999. Internet Society.
21. Wenjuan Xu, Mohamed Shehab, and Gail-Joon Ahn. Visualization based policy analysis: Case study in SELinux. In *Proceedings of the 13th ACM Symposium on Access Control Models and Technologies (SACMAT-2008)*, pages 165–174, June 11–13, 2008, Estes Park, Colorado, 2008.