

Purdue University

Purdue e-Pubs

Department of Computer Science Technical
Reports

Department of Computer Science

1991

Preventing Weak Password Choices

Eugene H. Spafford

Purdue University, spaf@cs.purdue.edu

Report Number:

91-028

Spafford, Eugene H., "Preventing Weak Password Choices" (1991). *Department of Computer Science Technical Reports*. Paper 875.

<https://docs.lib.purdue.edu/cstech/875>

This document has been made available through Purdue e-Pubs, a service of the Purdue University Libraries.
Please contact epubs@purdue.edu for additional information.

PREVENTING WEAK PASSWORD CHOICES

Eugene H. Spafford

CSD-TR-91-028

April 1991

Preventing Weak Password Choices*

Purdue Technical Report CSD-TR-91-028

Eugene H. Spafford
Department of Computer Sciences
Purdue University
West Lafayette, IN 47907-2004
spaf@cs.purdue.edu
(317) 494-7825

March 1, 1991

Abstract

A common problem with systems that use passwords for authentication results when users choose weak passwords. Weak passwords are passwords that are easy to guess, or likely to be found in a dictionary attack. Thus, the choice of weak passwords may lead to system compromise.

Methods exist to prevent users from selecting and using weak passwords. One common method is to compare user choices against a list of unacceptable words. The problem with this approach is the amount of space required to store even a modest-sized dictionary of prohibited password choices.

This paper describes a space-efficient method of storing a dictionary of words that are not allowed as password choices. Lookups in the dictionary are $O(1)$ (constant time) no matter how many words are in the dictionary. The mechanism described has other interesting features, a few of which are described here.

Keywords: passwords, dictionaries, password aging

*This has been submitted to the 14th National Computer Security Conference.

1 Introduction

Passwords are a commonly-used method of authentication. A unique sequence of characters is presented to the system when identification is needed. This sequence is then compared with a stored sequence, perhaps after some transformation (e.g., encryption). A match provides the proof of identity.

One weakness with password systems is the choice of the password. If the choice of possible characters to use in the password is too small, or if the overall length of the password is too short, the password may be compromisable. Even a rich character set may not be sufficient to create secure passwords if the combination of characters is restricted to an arbitrary set of possibilities. Thus, good password choice should avoid common words and names (cf. [15, 1, 10, 6]).

As an example, consider the UNIX¹ password system. This password mechanism is based on a cryptographic transformation of a fixed string of zero bits, using the user-supplied password as a key.[12] The transformation is an altered version of DES encryption, performed 25 times. The transformation is sufficiently slow so that exhaustive key-space attacks are currently not practical, although fast implementation such as *deszip*,[3] can perform many thousands or tens of thousands of comparisons per second.

In UNIX, the encrypted version of the password has traditionally been kept in a world-readable file; the safety of the passwords has been protected by the time complexity of an exhaustive attack. Thus, one of the keys to the safety of UNIX passwords is a large potential key-space for passwords. If the full character set is used, and seven or eight-character passwords are chosen, the number of potential passwords to be searched is far too large to be successfully searched, even at high speed.² Unfortunately, users often select passwords that do not exploit the large key-space available. Instead, they choose common words and names, or simple transformations of those names. This greatly simplifies an attacker's task.

This tendency to select weak passwords has led to a number of system break-ins, some quite highly publicized: [21, 20, 18, 23, 14]. Current technology is such that construction of a large pre-encrypted dictionary on-line using optical disks is easily done. By creating such a dictionary, a password

¹ UNIX is a trademark of Unix System Laboratories, Inc.

² Assuming a usable character set of 120 characters, there are approximately 4.3359×10^{16} possible passwords. At 10,000 attempts per second, an exhaustive search of this key-space would require over 137,000 years to complete.

search and attack may be easily conducted in a matter of seconds. Without such a database, but using a tool such as *deszip* on a modern workstation, it is possible to make a full scan of 300,000 dictionary entries against several hundred passwords in a matter of a few days.

Despite wide-spread publication of good password policy and the risks inherent in bad passwords, users continue to select weak passwords. This is a continuing threat to the best-managed systems. (For example: [9, 2, 15, 8, 7, 10, 6, 11, 19, 22, 24].)

There are four basic methods for a system administrator to enforce better password security on a computer system:

1. Educate and encourage users to make better choices of passwords.
2. Generate strong passwords for users and do not allow them to choose passwords of their own creation. This is often done using some random password generator.
3. Check passwords after-the-fact and force users to change those that can be easily broken with a dictionary attack.
4. Screen users' password choices and prevent weak ones from being installed.

This first method, that of educating users to choose strong passwords, is not likely to be of use in environments where there is a significant number of novices, or where turnover is high. Users might not understand the importance of choosing strong passwords, and novice users are not the best judges of what is "obvious." For instance, novice users (mistakenly) may believe that reversing a word, or capitalizing the last letter makes a password "strong."

A further problem is if the education provided to users on how to select a password is itself dangerous. For instance, if the education provided gives users a specific way to create passwords, such as using the first letters of a favorite phrase, then many of the users may all use that algorithm, thus making an attack easier.

The second method of strengthening passwords is to generate the passwords for the users and not allow them the opportunity to select a weak password. For this mechanism to work well the passwords need to be randomly drawn from the whole key-space. Unfortunately, this method also has flaws. In particular, the "random" mechanism chosen might not be truly

random, and could be analyzed by an attacker. Furthermore, random passwords are often difficult to memorize (especially if they are changed (*aged*) regularly). As a result, users may write the passwords down, thus providing an opportunity to intercept them without the effort of a dictionary search.

The third method of preventing poor password choice is to scan the passwords selected, after they are chosen, to see if any are weak. This is supported by many systems, including *deszip* and COPS.[5] There are significant problems with this approach:

- The dictionary used in the search may not be comprehensive enough to catch some weak passwords. Outside attackers might think of these choices, but the password scanner would not include them in the search.
- The scanning approach takes time, even for a fast implementation. A lucky (or determined) attacker may be able to penetrate a system through a weak password before it is discovered by the scanner. This is especially a problem in an environment with a very large number of users.
- The output of a scanner may be intercepted and used against the system.

Additionally, there is not always a correlation between finding a weak password and getting it replaced with a stronger one. At my university, for instance, at least two faculty members have repeatedly been informed of the weakness of their passwords as exposed by a scanner, but they have not chosen new passwords in over three years.

The fourth method, that of disallowing the choice of poor passwords in the first place, appears to have none of the drawbacks mentioned above. However, it too has difficulties associated with it. In particular, the storage required to keep a sufficiently large dictionary may prevent this method from being used on workstations and small computer systems. For instance, the standard UNIX dictionary, `/usr/dict/words`, is about 25,000 words and 200,000 bytes of space. A dictionary of 10 to 20 times that size would be necessary for reasonable protection; there are over 170,000 words in Webster's New World Dictionary, and that would occupy well over a million bytes of disk storage. That figure does not include many slang and colloquial words and phrases, nor does it include any user names, local names and phrases, likely words in foreign languages, or other strings shown to be poor pass-

word choices. A comprehensive dictionary might double these figures, and require over three million bytes of storage.

Maintaining such a large dictionary is also difficult. To add new words or phrases means that the dictionary must either have additional space overhead for indexing, it must be sorted after each addition, or else lookups are linear in the length of the dictionary. Again, in small computer environments, none of these alternatives may be appropriate.

2 OPUS

The OPUS Project³ is intended to address the space problems associated with a sufficiently complex password screening dictionary. The goal is to derive a mechanism that provides protection equivalent to a comparison against a large dictionary, yet be small enough to be practical in a small computer environment.

2.1 The Dictionary Filter

The main component of this system is a Bloom filter-encoded version of the dictionary.[4] A Bloom filter is a well-studied probabilistic membership checker, often used in applications such as spelling checkers.[16, 13, 17] It works as follows: a word to be entered into the filter is passed through n independent hash functions generating integer values. Each of these values is used as an index into the filter, represented as a bitmap. The bits corresponding to the input word are then set. This procedure is repeated for each word to be entered into the filter.

When a lookup is to be performed, the word to be examined is passed through the same hash functions and the corresponding bits in the filter are examined. If any of the bits is reset (i.e., not set), then the word is determined not to be present in the dictionary. If all the corresponding bits are set, the likelihood is high that the word was in the list that was used to build the dictionary, and in the case of OPUS, rejected as a weak password choice. The probability of a false rejection can be set arbitrarily low by increasing the size of the bitmap and increasing the number of hash functions used; an obvious upper bound is the size of the plaintext dictionary.

To be more exact, assume we have a hash table of N bits, and d non-overlapping hash functions. From [4], with n words we have the proportion

³ Obvious Password Utility System.

of bits left unset, ϕ , equal to

$$\phi = \left(1 - \frac{d}{N}\right)^n$$

A word will be falsely shown as present in the dictionary iff it hashes to a set of bits that are all set. The expected proportion, P , of words that will be mistakenly shown as in the dictionary is thus

$$P = (1 - \phi)^d$$

From these equations, we can derive appropriate values to choose for our filter and hash functions.

For example, suppose we pick $n = 250,000$ words for the dictionary, and we wish to have a 0.5% chance ($P = 0.005$) of false positives on any given text string. If we choose six uniform hash functions, we will need 2,800,000 bits of storage and achieve $\phi = 0.586$. This works out to a file of 350K bytes. Reducing $P = 0.01$ results in needing only 300K bytes of storage for the dictionary. Storing the full dictionary as plaintext would likely take in excess of 2 Mb of storage. Thus, we are able to achieve almost a seven-fold compression with only a small loss of accuracy.

As can be seen from the above examples, with the appropriate choice of hash functions it is possible to greatly reduce the storage necessary to keep an extensive dictionary of words to compare against password choices. By making queries on the dictionary with variations of the candidate password — upper/lower case, reversed, trailing digit, etc. — it should be possible to quickly check for the strength of the password. Each probe into the dictionary is basically a constant-time operation, so the number of words in the dictionary has no effect on the time of access. If the union of all the probes results in a positive response, the user is told to try again.

2.2 Other Features

The model of the dictionary used in OPUS provides benefits other than simple dictionary lookup. By providing a writable interface to the dictionary for the system administrator, it is a simple task to add the representation of new words to the dictionary. The administrator can therefore augment the dictionary with local user names and colloquialisms. Adding words to the dictionary requires no expensive sort or temporary storage. Furthermore, the system administrator never needs to be concerned if a word has already been added — adding a word more than once has no effect.

The OPUS system also supports password aging. With password aging, users are required to change their passwords periodically. However, a common fault with password aging is that users attempt to reuse old passwords, and this may present a security risk.

OPUS can be configured so that whenever a password is changed, it is added to the dictionary. Thus, if a user attempts to reuse an old password, she will find it already in the dictionary, and the choice will not be allowed. As seen from the value of ϕ , above, there is plenty of room in the dictionary for adding new words, so even prolonged operation will not result in a noticeable degradation of service.

One obvious problem with updating the dictionary in this manner is the possibility of an attacker using delta information to craft a set of password attempts. That is, by observing the changes made to the filter when another user changes his password, an attacker might be able to use the hash functions to derive a set of possible text strings that account for the changes, and use these in a penetration attempt.

A related problem is if an attacker finds a way to use the dictionary as a filtering mechanism to exclude patterns when doing a brute-force key-space search to break passwords. Doing a probe into the dictionary will determine if a candidate is a possible choice or not, thus saving (some) on the computation required to perform an exhaustive search.

Luckily, there is a simple way to defeat these problems. Instead of hashing plaintext words into the dictionary, OPUS first encrypts the words to be entered or examined. The encryption will be something time-consuming, similar to multiple rounds of the DES function, and computationally infeasible to reverse. The hashing algorithms are then applied to the encrypted string rather than to the plaintext. Thus, to gain any information from the dictionary, either as a pre-screen or as a source of delta information, would require much more computational effort than some other approach (e.g., exhaustive key-space search).

3 Final Remarks

This paper has discussed the motivations and design behind a system for preventing users from installing weak passwords. The system should be compact and simple to customize and enhance. It can be used standalone, as a frontend to an existing password program, or coupled with some form of password generator so as to prevent the accidental generation of a word

susceptible to dictionary attacks.

The choice of hashing algorithms used with the system is critical for the success of the filter. Choosing non-uniform or overlapping hash algorithms reduces the effectiveness of the Bloom filter by increasing the incidence of false positives (effectively shrinking the number of useful bits employed). If possible, the hash algorithms should be chosen to produce the same results whether used on a string or on its reverse. This will allow probes for common words and their reverses to be made simultaneously. Case-insensitivity can also be used in the hash functions, but this may result in too great a narrowing of the key space; words in monospace, or with only a leading or trailing capital letter are perhaps the only combinations that need to be examined.

A UNIX version of OPUS is being constructed. It will be preloaded with a locally-developed dictionary of almost 300,000 words. Experiments will then be conducted to determine, for this dictionary, the optimal working size and number of hash functions. Further experiments will determine the accuracy rate for rejection of candidate passwords that are not present in the real dictionary, and the speed of operation. By performing side-by-side experiments with users selecting potential passwords and comparing a dictionary search with the results of the Bloom filter, it should be possible to determine the overall utility of this approach. The results from those experiments will be reported later.

References

- [1] Ana Maria De Alvaré. How crackers crack passwords, or what passwords to avoid. Technical Report UCID-21515, Lawrence Livermore National Laboratory, 1988.
- [2] Ana Maria De Alvaré and Jr. E. Eugene Schultz. A framework for password selection. Technical Report UCRL-99382, Lawrence Livermore National Laboratory, 1988.
- [3] M. Bishop. An application of a fast data encryption standard implementation. *Computing Systems*, 1(3):221-254, 1988.
- [4] Burton H. Bloom. Space/time trade-offs in hash coding with allowable errors. *Communications of the ACM*, 13(7):422-426, July 1970.

- [5] Daniel Farmer and Eugene H. Spafford. The COPS security checker system. In *Proceedings of the Summer Usenix Conference*. Usenix Association, June 1990.
- [6] Simson Garfinkel and Eugene H. Spafford. *Practical Unix Security*. O'Reilly and Associates, Inc., May 1991.
- [7] David L. Jobusch and Arthur E. Oldehoeft. A survey of password mechanisms: Weaknesses and potential improvements. part 2. *Computers & Security*, 8(8):675-689, 1989.
- [8] David L. Jobusch and Arthur E. Oldehoeft. A survey of password mechanisms: Weaknesses and potential improvements. part 1. *Computers & Security*, 8(7):587-603, 1989.
- [9] Daniel V. Klein. A survey of, and improvements to, password security. In *UNIX Security Workshop II*, pages 5-14. The Usenix Association, August 1990.
- [10] Belden Menkus. Understanding password compromise. *Computers & Security*, 7(6):549-552, December 1988.
- [11] Chris Mitchell and Michael Walker. The password predictor — a training aid for raising security awareness. *Computers & Security*, 7(5):475-481, October 1988.
- [12] Robert Morris and Ken Thompson. Password security: a case history. In *Unix Programmer's Supplementary Documentation*. AT&T, November 1979.
- [13] James K. Mullin. A second look at Bloom filters. *Communications of the ACM*, 26(8):570-571, August 1983.
- [14] Neil Munro. Simple password opens navy computer to hacker. *Government Computer News*, 7(15):61, July 1988.
- [15] National Computer Security Center. Password management guideline. Technical Report CSC-STD-002-85, US Department of Defense, 1985.
- [16] Robert Nix. Experience with a space efficient way to store a dictionary. *Communications of the ACM*, 24(5):297-298, May 1981.

- [17] M. V. Ramakrishna. Practical performance of Bloom filters and parallel free-text searching. *Communications of the ACM*, 32(10):1237-1239, October 1989.
- [18] Brian Reid. Reflections on some recent computer break-ins. *Communications of the ACM*, 30(2):103-105, February 1987.
- [19] Bruce L. Riddle, Muray S. Miron, and Judith A. Semo. Passwords in use in a university timesharing environment. *Computers & Security*, 8(7):569-578, 1989.
- [20] Donn Seeley. Password cracking: A game of wits. *Communications of the ACM*, 32(6):700-703, June 1989. 1989.
- [21] Eugene H. Spafford. The Internet Worm: Crisis and aftermath. *Communications of the ACM*, 32(6):678-687, June 1986.
- [22] Cliff Stoll. How secure are computers in the U.S.A.? an analysis of a series of attacks on milnet computers. *Computers & Security*, 7(6):543-547, 1988.
- [23] Cliff Stoll. *The Cuckoo's Egg*. Doubleday, NY, NY, October 1989.
- [24] Patrick H. Wood and Stephen G. Kochan. *Unix System Security*. Hayden Book Company, 1987.