Password Exhaustion: Predicting the End of Password Usefulness

Luke St.Clair, Lisa Johansen, William Enck, Matthew Pirretti, Patrick Traynor, Patrick McDaniel, and Trent Jaeger

Systems and Internet Infrastructure Security Laboratory
The Pennsylvania State University, University Park PA 16802
{lstclair,johansen,enck,pirretti,traynor,mcdaniel, jaeger}@cse.psu.edu

Abstract. Passwords are currently the dominant authentication mechanism in computing systems. However, users are unwilling or unable to retain passwords with a large amount of entropy. This reality is exacerbated by the increasing ability of systems to mount offline attacks. In this paper, we evaluate the degree to which the previous statements are true and attempt to ascertain the point at which passwords are no longer sufficient to securely mediate authentication. In order to demonstrate this, we develop an analytical model for computation to understand the time required to recover random passwords. Further, an empirical study suggests the situation is much worse. In fact, we found that past systems vulnerable to offline attacks will be obsolete in 5-15 years, and our study suggests that a large number of these systems are already obsolete. We conclude that we must discard or fundamentally change these systems, and to that effect, we suggest a number of ways to prevent offline attacks.

1 Introduction

Password-based authentication mechanisms are the primary means by which users gain legitimate access to computing systems. Because of their central role in the protection of these systems, the vulnerabilities inherent to these methods have long been known throughout the security community. The best known of these vulnerabilities is password choice. A variety of studies [21,25,32] cite the lack of *entropy*, or unpredictability, included in each password as the root of the problem. Because of the chronic under-use of the available key space, as many as 30% of user passwords are recoverable within a period of hours [24].

The common wisdom is that if users can be educated to select "perfect" passwords, offline brute-force attacks to recover such information will remain beyond the computational ability of modern machines [19]. In reality, the current entropy in a perfectly-random 8 character password, the most common password length, is actually less than that of a 56-bit DES key[1]. Thus, the security provided by these passwords is questionable. In order to increase the security provided by passwords, password length increases and password policies are commonly employed. A variety of password policies now request 15 character passwords. In this case, the entropy is comparable to 3DES or AES.

[1] DES was effectively broken by a brute-force attack in 1999 [2].

A. Bagchi and V. Atluri (Eds.): ICISS 2006, LNCS 4332, pp. 37–55, 2006.
© Springer-Verlag Berlin Heidelberg 2006

Kahoot! ASA et al. Ex. 1024-001

Password policies for guiding users to select more effective passwords have become more prevalent. As systems continue to rely on passwords to provide authentication security, it is important to investigate the validity of these improvements.

In addition to the future increases in computing power, the viability of password systems is limited by the entropy that humans are actually able to use in practice. Given than human beings are only capable of remembering approximately seven random items [10], an increase in password length does not necessarily mean a commensurate increase in real entropy. As passwords lengths increase, users may develop techniques to use predictable chunks of randomly arranged passwords. Also, users will be tremendously challenged to memorize multiple passwords of such length.

In this paper, we investigate two fundamental claims: (1) near-term increases in available computing power will soon enable offline brute-force cracking of perfectly-random 8 character passwords on a variety of commodity platforms, and (2) the maximum entropy that we can expect from a password is limited to no more than the commonly used 8 characters we have already, thus rendering password systems that permit offline attacks obsolete. First, we use current forecasting of hardware performance to estimate the end of the computational infeasibility for offline password attacks given the entropy of an 8 character password. We find that computing power that should be easily available to a typical user will be sufficient to break a perfectly random 8 character password in May 2016. For more motivated attackers, the time to recover passwords will be insignificant. Secondly, we examine the entropy of real passwords and the impact of password policies upon that entropy. NIST has done an analysis of the entropy present in real-world passwords, and based on this measure and the capabilities of modern password cracking tools, an attacker with only one machine who could search the potential password space in order of increasing randomness would be able to recover even an 8-character password including numbers and symbols in less than 15 hours. We find that in real passwords of the CSE department at Penn State, the entropy is only slightly better than this figure. Further, we find that password policies significantly limit password entropy and do not appreciably improve the protection of passwords. No solution is known to exist that can save password systems susceptible to offline attacks from obsolescence in the near future.

The remainder of this paper is organized as follows: in Section 2, we discuss predictions of future computer performance and their bearing on password vulnerabilities; Section 3 examines the ways in which entropy is actually removed from systems and revisits the above predictions; Section 4 considers solutions to this problem; related works are presented in Section 5; Section 6 offers concluding remarks.

2 Future of Password Recovery Power

This section considers how hardware improvements and processor availability impact the security provided by password authentication systems. We begin by introducing a model of computing used to assess the vulnerability of password systems to offline attacks. Using this model, we consider the present and future security of popular password systems.

Kahoot! ASA et al. Ex. 1024-002

2.1 Forecasting Model for Password Recovery

To assess the viability of current and future password systems, we introduce a model for investigating the impact of increasing processor speeds and parallelism on brute-force attacks. These factors, modeled as functions $s(t)$ and $p(t)$, are based on expert predictions of future computing trends. We also evaluate the effect of the growing availability of large systems of computers on brute-force attack speed.

Model Definition

Our model is composed of the following components: password space, processor performance, parallelism, and system size.

Password Space(c): The password space is the set of all possible passwords that a system can represent. In terms of password recovery, the password space indicates the average amount of work required to recover a password. Given the limitations of human memory, we shall assume a typical user password is composed of 8 characters, where each character can be any of the 95 characters readily represented with a keyboard. In the best case scenario (from the point of view of system security), user passwords will be uniformly distributed across the password space. Thus, an adversary on average must search half of the password space to recover a password. Based on these parameters, we represent the average number of tries required to recover a password as a constant:

$$c = 95^8/2 \approx 3.3 \times 10^{15}, \tag{1}$$

where each attempt to break a password is termed as *try*.

Processor Performance($s(t)$): The processor performance function models the amount of work that can be accomplished by a single processing element[2]. To map this factor to password recovery, we define processor performance as the number of seconds required to perform a single try, denoted as a time varying function $s(t)$. In Section 2.1 we consider several models that predict how processor performance will change with time.

Parallelism($p(t)$): The parallelism function models the increasing prevalence of processor replication in contemporary computing systems by measuring the number of processor cores present within a single computer. Password recovery is a highly parallelizable activity; the password space can be subdivided into disjoint components and independently processed by different processor cores. The level of parallelism present in a given machine greatly increases the rate at which the password space is examined. For instance, a machine with 4 processing cores can simultaneously perform 4 tries. We denote the level of parallelism present in a given computer as a time varying function $p(t)$. In Section 2.1 we consider different models that have been used to forecast the number of independent processing cores present within a single computer.

System Size(z): System size models the increasing prevalence of computational devices. For instance, the number of computers in homes is steadily increasing [16]. Further, the number of computers present in computing clusters is quickly growing. Finally, the overwhelming size of botnets is increasing. To capture this trend we represent the number of independent computers present in a system as a variable z.

[2] To avoid ambiguity, this factor specifically measures the amount of work a single processor core can perform.

Kahoot! ASA et al. Ex. 1024-003

Password Cracking Forecasting Model($T(t)$): Given our definitions of $c, s(t), p(t)$, and z we can now introduce our model of forecasting how the computing trends of increasing processor performance and increasing parallelism will affect the viability of brute-force password cracking attacks. Our model represents the amount of time required to recover a random 8 character password:

$$T(t) = \frac{(3.3 \times 10^{15}) \cdot s(t)}{p(t) \cdot z}. \tag{2}$$

Predicting Processor Performance $s(t)$

This section examines predictions on the growth of future processor performance as made by experts in the field. It then defines the function for this growth, $s(t)$, over time which is used in our model.

Determining future processor performance has been a widely studied problem for more than 50 years. Moore stated in 1965 that chip density will double every 12 to 18 months. Unfortunately, chip density is reaching its limits due to heat and power consumption [14]. Because of these challenges, the industry is looking towards other methods to increase overall computing power instead of focusing on clock speed. Nanotechnologies, compiler optimization and other innovations are being considered as approaches for increasing chip performance [14,20]. However, the industry still looks to Moore's law as a predictor of future computing power [20].

Shown in Figure 1(a), Moore's Law is represented by the function, $s_M(t)$. However, studies have shown that the rate of performance growth proposed in Moore's law is unrealistically fast [18]. One of these studies states that, over the past 7.5 years, the actual rate of computer performance growth has been closer to 41% per year. This more conservative predictor serves as a second function for performance growth, $s_R(t)$.

Because the limits of physics are affecting the application of traditional methods for performance improvement, $s_M(t)$ and $s_R(t)$ may be unrealistic for future predictions. Because chip density reaching its limit, there has been much discussion that Moore's law is no longer valid [17]. Experts are beginning to doubt that processor power is going to continue to grow at rates that we have seen in the past. Taking this into consideration, we define two more functions, $s_{SG}(t)$ and $s_0(t)$, to more conservatively project the growth of processor performance. The first function, $s_{SG}(t)$, is a conservative estimate of 20% processor power growth per year derived from the expert predictions of future computing power. The second function, $s_0(t)$, assumes no growth over the next 15 years. All four of our processor performance growth functions are plotted in Figure 1(a).

Predicting Parallelism Factor $p(t)$

Parallel computing is popularly seen as a counterbalance to slowing growth of processor performance. Multi-core technologies, multiple processors and hyper-threading have been proposed to increase performance. Intel and AMD have released estimates for the amount of parallelism that they expect will exist in a single computer in the near future. Intel believes that a processor will have anywhere from tens to hundreds of cores within ten years [14]. AMD projects that processors will contain more than 2 cores by 2007 and more than 8 by 2008 [1].

Kahoot! ASA et al. Ex. 1024-004

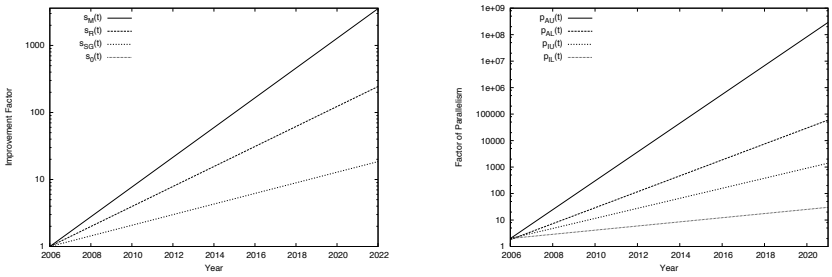

(a) Processor Performance (b) Parallelism Impact

Fig. 1. Future Computing Performance Estimates

Given these estimates, we predict where parallelism will be in the near future. We extrapolate two functions for AMD's estimates and two functions for Intel's estimates. The first function of parallelism growth, $p_{AU}(t)$, is based on AMD's upper estimate of the number of processors available in a single computer while $p_{AL}(t)$ is their lower estimate. Similarly, Intel's upper estimate defines the function $p_{IU}(t)$ and their lower estimate defines $p_{IL}(t)$. From the graph of these four functions shown in Figure 1(b) we see that AMD's estimates, $p_{AU}(t)$ and $p_{AL}(t)$, are the upper approximations for parallelism growth while Intel's estimates, $p_{IU}(t)$ and $p_{IL}(t)$, establish the lower bound.

System Size z

Another way to gain parallelism in the password recovery computations is to use multiple computers. The more computers that are used, the less time it takes to recover a password. A system of computers can consist of personal computers, clustered nodes, or large networks like botnets. Any personal computer user can own any number of computers. In the case of a computing cluster, the size of these systems can be much larger. Sizes of 20 to 500 nodes are not atypical for today's standard computer clusters. Large networks of computers provide the largest factor of parallelism when executing one task. A botnet, a large number of compromised machines that can be remotely controlled by an attacker, is an example of such a network. Botnets can range in size from thousands to hundreds of thousands of computers. This growing availability of computers directly affects the amount of parallelism available to any user.

2.2 Future Password Recovery

This section depicts the extent to which computing performance improvements threaten the security of password authentication systems as determined by our model. We begin by briefly describing the password systems that we attempt to break and the time to recover passwords for each system on today's commodity hardware. The processor performance, $s(t)$, parallelism factor, $p(t)$, and system size, z, of the future computing systems used in the experiment are defined. The ability to recover passwords of these future systems is then evaluated in Figure 2 and discussed.

Kahoot! ASA et al. Ex. 1024-005

In our experiments, we analyze three password systems: Unix crypt, MD5 crypt and Kerberos. Unix crypt [9] is based on the DES algorithm and, with the increasing processing power available, has become vulnerable to brute force attacks. The MD5 algorithm [30] is also used in Unix-based systems and has been implemented as an upgrade from the DES algorithm. Compared to DES's 56-bit key size, MD5 uses a 128-bit value making brute force attacks comparatively more difficult. Kerberos [23] is a widely-used single-sign-on network authentication system. Tickets in the Kerberos system are encrypted with a key derived from a user's password. These tickets can be attacked in order to recover that password. As explained in the introduction, we chose to analyze these systems due to their current wide spread use. For more information about these password systems see Appendix A. We ran password recovery software on commodity hardware to determine the speed of tries for each password system. The results from these experiments are illustrated in Table 1. These values serve as indicators of current day password recovery ability. Using them as inputs to our models, we can derive the password recovery ability of future computing systems.

Table 1. Password Recovery Speeds

System	Tries/Sec
Unix Crypt	1557890
MD5 Crypt	17968
Kerberos	55462

For the following analysis, we posit one possible future computer type. Given the fact that chip technologies are reaching the limits of power and heat, the slow growth processor performance function, $s_{SG}(t)$, is used. In order to determine a median parallelism factor, we take the average of the two middle values; AMD's lower estimate and Intel's higher estimate. This results in a function, $p_{AVG}(t)$, that characterizes the average of $p_{AL}(t)$ and $p_{IU}(t)$. We analyze the impact of this computing power as it exists within these systems of the future with the following model:

$$T_z(t) = \frac{(3.3 \times 10^{15}) \cdot s_{SG}(t)}{p_{AVG}(t) \cdot z}. \tag{3}$$

The number of computers within a system, the parameter z in Equation 3, can range anywhere from one to hundreds of thousands. We examine a six different systems in our study: a personal computer system consisting of 2 computers, clusters of 10 and 100 nodes, and botnets of 1000, 10000, and 100000 compromised hosts.

Beginning with the initial values presented in Table 1, we were able to determine, for each password system, what each of the modeled computer systems was capable of recovering over the span of 15 years. The results from these experiments are presented in Figure 2.

The most apparent result is that a botnet with 10,000 or more compromised computers is currently able to recover any password from any password system in under 6 months. In less than five years, any botnet with at least 1,000 compromised computers can recover any password in under a month.

Kahoot! ASA et al. Ex. 1024-006

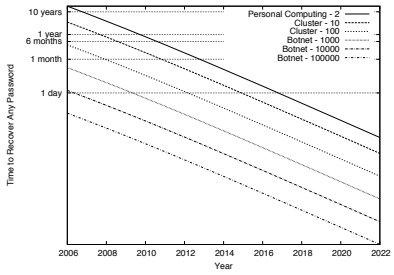

(a) Unix crypt

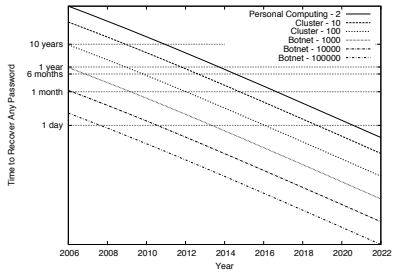

(b) Kerberos

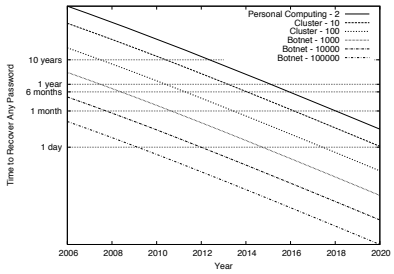

(c) MD5 crypt

Fig. 2. Future Password Recovery Analysis

Given a smaller system, like a cluster, we see that password recovery, naturally, takes longer. An average sized cluster is able to recover any Unix crypt password in under 6 months today. However, within only 8 years, a cluster of minimal size will be able to recover any password from our three presented password systems.

Examining the extreme case of personal computing systems, the results are startling. Within three years, any Unix crypt password will be recoverable in under 6 months. Unix crypt is obviously broken. The more devastating result is that any password from any of the other evaluated systems is recoverable in under 6 months by a single personal computing system in 10 years. This means that our trusted authentication systems will be vulnerable to raw computing power from the comfort of your own home before 2017.

3 Passwords in Practice

In this section we show that the current state of password security is actually much worse than the theoretical model presented in Section 2.1 suggests. The preceding model examines the effect of improving hardware on password recovery, but does so considering the full password space of 95^8 possibilities. In reality, the password space is often much smaller, thus an adversary is not required to try every one of the 95^8 possible passwords. For example, password policies serve to reduce the amount of work an attacker is required to do to recover random passwords by reducing the possible

Kahoot! ASA et al. Ex. 1024-007

password space. We examine specific password policies and demonstrate the effect that they have on the speed of brute force attacks.

Password systems are further weakened due to the poor choice of passwords. In practice, users choose non-random passwords that contain much less than their maximum allowable entropy. We demonstrate the degree to which this occurs by examining passwords within an actual institution and also by discussing NIST's study of passwords and their actual entropy. The effect of this reduced entropy on the speed of password recovery is examined.

3.1 Password Policy Restrictions

Based on the recommendations of the security community [8,31], many password systems have begun to implement password policies restricting the types of passwords that may be chosen. For instance, some sites do not allow users to choose characters outside the alphanumeric set. Others require passwords to be between a minimum and maximum length. Still others make restrictions on the types of characters that must be present in a password. While these rules help to prevent users from choosing dictionary-based passwords, we show that they decrease the total password space and that this is not effective in preventing brute-force attacks. We examine a number of policies and evaluate their effect on the speed of password recovery attacks in comparison with the results presented in Section 2.1.

We first analyze how policy restrictions reduce the number of possible passwords. Let R_i be the set of passwords that do not satisfy a certain policy i. For instance, if policy i required users to choose a lower-case letter $|R_i| = (95-26)^8$ thus, the password space is $|\neg R_i| = 95^8 - (95-26)^8$. We also define $R_i \cap R_j$ to be the intersection of passwords that do not satisfy both policies i and j (i.e. both policies are not satisfied). Now, we apply a variant of the inclusion-exclusion principle[3] to the total password space to get the following formula, which computes how many passwords satisfy all of the policies specified:

$$|\bigcap_{1\leq i\leq k} \neg R_i| = (95^8)/2 + (-1)^1(\sum_{1\leq i\leq k} |R_i|) + (-1)^2(\sum_{1\leq i_1\leq i_2\leq k} |R_{i_1} \cap R_{i_2}|) + \\ \ldots + (-1)^{k-1}(|R_1 \cap R_2 \cap \ldots \cap R_{k-1} \cap R_k|) \quad (4)$$

The first password policy we examine is from a recommendation made by the SANS Institute password policy page [8]. SANS (SysAdmin, Audit, Network, Security) is a large collaborative group of security professionals that provide information security training and certification [7]. Their recommended policy is intended to be used by businesses when they establish password policies for their enterprise networks. SANS recommends that users pick at least one upper and lower case character, 1 digit and 1 special character. Applying the formula in Equation 4, this fairly typical policy reduces the number of valid passwords by more than a factor of 2.

[3] The inclusion-exclusion principle is used to determine the cardinality of multiple finite sets without double counting. It over compensates by repeatedly including set intersections, then recompensates by excluding the excess intersections.

Kahoot! ASA et al. Ex. 1024-008

The Computer Science and Engineering department at The Pennsylvania State University recently enacted a password policy applying to the password choices of all students and faculty in the department. Following the common wisdom, they used the Sun password policy mechanisms to define a policy requiring all users to have 2 upper case characters, 2 lower case characters, 1 digit, and one special character. Using the previous formula, this reduces the pool of potential passwords by nearly a third.

The last set of potential passwords examined here are those generated with the pwgen utility [6]. pwgen is a Unix utility which generates "memorable" passwords of user-defined size (default is 8 characters). However, until recently, pwgen would only mix alphanumeric characters randomly to form passwords, and would not use symbols. According to the pwgen changelog [5], this was done so that passwords would be "much more usable." Obviously, this greatly restricts the number of potential passwords that can be chosen to the size 62^8, down from 95^8. Unfortunately, pwgen is widely used to generate "random" passwords when secure initial or replacement passwords are needed.

We now examine how limiting the password space affects the speed of brute-force attacks. We perform this examination with the same models used in Section 2. Because the previous section already demonstrated that older Unix crypt hashes are too weak to provide practical security, we evaluate the speed of brute-force attacks under the Kerberos system. In this way, we demonstrate the degree to which attacks can be sped up in a system that would otherwise remain somewhat secure for the next few years.

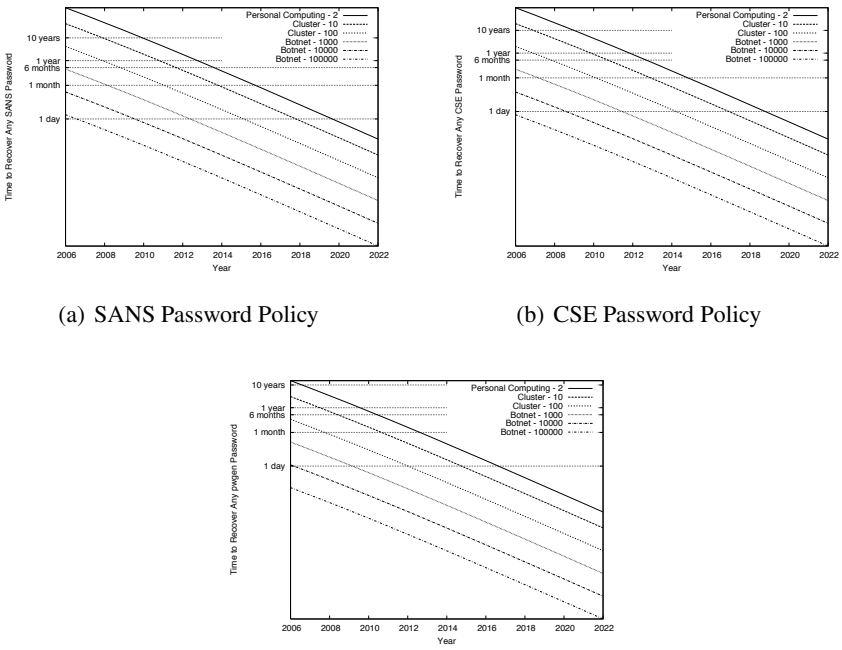

(a) SANS Password Policy

(b) CSE Password Policy

(c) pwgen Passwords

Fig. 3. Password Recovery with Policy Restrictions

Kahoot! ASA et al. Ex. 1024-009

The new estimates for future brute-force attacks are shown in Figure 3. As a point of reference, under the full password space, a single user will be able to crack a Kerberos password in under a year in October, 2013. We see that, under the restricted password space that SANS defines, this will happen around October, 2012. The Computer Science and Engineering policy restricts that password space such that a personal computer will be able to recover a Kerberos password in one year by around November 2011. And, most devastatingly, pwgen-generated passwords can be recovered approximately 30 times faster than passwords without restrictions. This makes pwgen passwords approximately as weak as older Unix crypt hashes without password restrictions. This policy, the most restrictive of the three, creates a situation in which a single personal computer can crack a Kerberos password in 1 year around June, 2009. This demonstrates how policy decisions can negate benefits accrued through proper algorithm choice

We conclude that password policies, while useful for eliminating the weakest passwords, can severely restrict the pool of passwords attackers must search. Requiring users to pick only a subset of potential passwords can drastically reduce the password space provided by perfectly random passwords in that system. The time to crack a perfect password is reduced by up to a factor of 30, as shown in Figure 3(c). It is important to note that when password policies are enforced, this sizable restriction is applied to each and every password, no matter how random the user's password is.

3.2 User Passwords

Despite password policies which try to force users to create more random passwords, users still choose passwords which contain very little entropy. Users often pick words with obvious letter replacements, like "0" for O, dictionary words with numbers or characters appended, and misspellings of common words. Also, by nature, users are much more likely to pick certain characters than others based on elements of a language [27]. In short, they still pick passwords that are not truly random and thus are vulnerable to intelligent password guessing attacks. In response to these tendencies, most password recovery tools today contain fairly sophisticated methods of guessing words with common symbol-for-letter replacements, words with numbers appended, and variations on words out of the dictionary. Some password crackers also have an intelligent brute-force mode which tries all possible passwords, but in order of increasing likelihood. For instance, the string "bgtyae1T" would be tried long before "t,I}&[*v". Then end result of this is that the brute-force times depicted in Figures 2 and 3 are still much too conservative when applied to actual passwords.

In order to evaluate the severity of weakly chosen passwords, a password file for the entire Computer Science and Engineering Department at Penn State University was obtained as described in Appendix A. This recovered the password hashes for 3500 users. The password recovery jobs were then submitted to a cluster of 20 nodes with dual AMD Opteron 250 processors. 16 of the 20 nodes were used for brute-forcing passwords based on character frequencies, and 4 nodes for trying passwords derived from dictionary words. The password recovery tool John the Ripper was used since it is widely available and has good support for both dictionary and brute-force based attacks. This program ran for 5 days, with the number of passwords recovered as a function of time graphed in Figure 4.

Kahoot! ASA et al. Ex. 1024-010

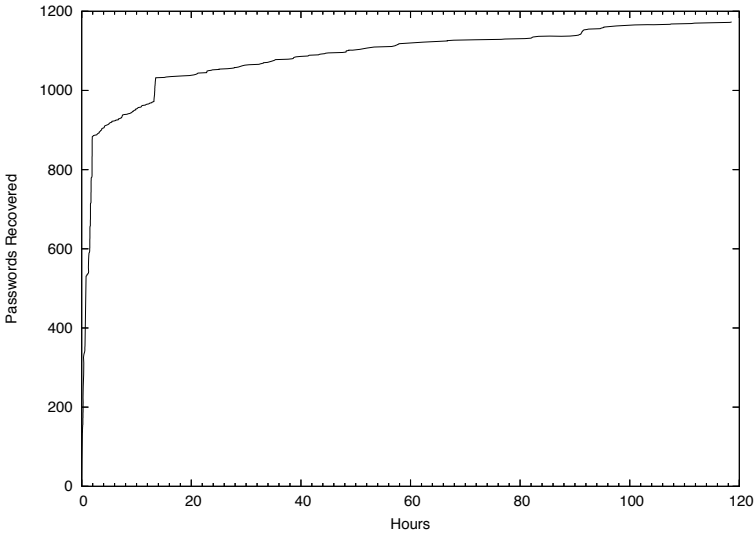

Fig. 4. Time to Recover Penn State CSE Passwords

One of the most startling observations from this graph is that approximately 25% of all passwords were cracked within the first 2 hours. It is particularly interesting to note that only 10.1% of the passwords recovered were recovered as a result of guesses based on dictionary words. The remaining 1118 passwords were all recovered by the nodes performing an intelligent brute-force attack. The brute-forcing method found in John the Ripper [3], which we used for these experiments, is clearly able to recover a large number of non-dictionary based passwords. This is a somewhat surprising result, as common lore in the security community is that the biggest problem with passwords is that users choose commonly used or dictionary-based passwords. Instead, this data may reflect both a growing consciousness about the weakness of dictionary passwords and a persistent inability of users to pick passwords uniformly from the potential space of all possible passwords.

3.3 Study of User Passwords

Although the results in the preceding section paint a dire picture of password strength in practice, there is evidence that the actual resistance of passwords to attack is much worse. According to the National Institute of Standards and Technology (NIST), passwords in practice are much weaker than the theoretical maximum previously discussed [32]. Based on experimental findings, NIST has given some guidelines as to the practical entropy provided by passwords that users choose. In this section, we evaluate the time to crack an 8 character password using password entropy guidelines from NIST.

NIST measures the amount of randomness within a given password using bits of entropy. Each bit of entropy under a given password policy multiplies the possible password space by 2 (e.g., 5 bits of entropy means there are 2^5 possible passwords). In a

Kahoot! ASA et al. Ex. 1024-011

system where a user may chose any 8 character password, NIST specifies that the first character of a password gives 4 bits of entropy, and successive characters give 2 bits apiece. We will call this the *permissive policy*. In many systems, however, users are forced to choose both upper-case and non-alphanumeric characters. This nets a password with an additional 6 bits of entropy, according to NIST. We will refer to this as the *special-character policy*. If passwords are checked against an exhaustive dictionary of at least 50,000 words (a factor of 10 less than the dictionary searched in the preceding results), 6 bits of entropy are added. Two new policies can be constructed, the *dictionary policy* and the *special-character and dictionary policy*, when the dictionary is applied to the permissive and special-character policies, respectively.

Given the previously discussed entropy values and equation for password space as determined by these values, we estimate the amount of time required to recover any password with a single machine and with a 20-node cluster. Table 2 presents the amount of time required to recover these passwords in theory as determined by NIST. Note that all of these measurements assume full 8 character passwords, as the time to compromise with 7 or fewer characters became negligible.

Table 2. Time to recover passwords as specified by NIST

Password Type	Bits of Entropy	Time to recover for single machine	20-node Cluster
permissive	18	$<$ 1 minute	$<$ 1 second
special-character	24	14 minutes	20 seconds
dictionary	24	14 minutes	20 seconds
special-character and dictionary	30	15 hours	22 minutes

Using the NIST entropy guidelines, our estimates predict faster recovery than our experimental results show, though not by much. As Table 2 shows, the average special-character and dictionary password, the strongest password type that the government identifies, takes 15 hours to recover. However, only a third of the CSE user passwords were recoverable in this time. The discrepancy in these values is likely caused by better passwords chosen within the computer science department as opposed to the general population.

Software may also play a role in explaining the discrepancy between the actual brute-force recovery speed and the theoretical brute-force recovery speed. For example, although John the Ripper is a sophisticated program, it does not do a perfect job of guessing passwords in order of increasing randomness. Thus, even though a user password may not be very random, it may not be quickly recovered due to the imperfect nature of the password recovery software. Our experience has confirmed this, as we have seen many passwords (like “myPword!”) take a long time to crack, while possessing little randomness. Software algorithms for guessing passwords have room for improvement, and the recent development of new password recovery methods [3] indicates that these improvements will likely continue to be made. This would decrease the time to recover a password on every system in every configuration for every password.

Kahoot! ASA et al. Ex. 1024-012

4 Mitigating Password Vulnerability

This section considers ways of mitigating the vulnerabilities of current password systems, now and in the future. We consider two broad approaches to limiting the vulnerabilities associated with passwords: the first is to simply prevent offline attacks from occurring, and the second is to reduce the effectiveness of the offline attack.

4.1 Preventing Offline Attacks

In order to mount an offline password attack, recovery material must be obtained, for example, a password file or a TGT. This material normally consists of encrypted or digested versions of passwords. Password material can be obtained actively or passively. Active recovery requires the adversary to perform some observable behavior such as initiating a fake login or reading a password file in order to obtain the necessary material. Passive recovery is a covert action such as eavesdropping a network exchange to acquire the material. Preventing the adversary from obtaining this password material would effectively prevent offline attacks. The mechanism used to prevent the attack is related to the means by which the password material is obtained.

In an attempt to prevent active password material recovery, recent UNIX systems have begun to provide mechanisms to reduce the visibility of password material to all users of the system. Recent versions of Unix introduced the concept of the `/etc/shadow` file, which stores the password hash in a file readable only by root, instead of by all users. This makes active recovery schemes more difficult, thus making offline attacks more difficult.

Preventing passive recovery of password material involves eliminating the availability of the material on unprotected networks. One option is to not send password material over the network. Password material includes anything that is encrypted with a password or a derivation of a password. Thus, in order to abstain from sending password material, data must be encrypted with something besides the password. The secure remote password protocol (SRP) [33] and similar protocols [11] have created a way for two parties to agree on a symmetric session key with which to encrypt data instead of using passwords. However, authentication is still performed with passwords, as both parties must have knowledge of the password in order to agree on the symmetric key. Thus, these protocols have eliminated the need to send password material over an insecure network in order to support authentication. They are designed such that they effectively prevent brute-force online and offline attacks.

Protecting the network over which passwords are sent is another way to protect password material. This technique is useful in systems that cannot support changes to their protocols. For instance, SRP could be difficult to apply to authentication to financial web services, due to time synchronization restraints, export restrictions, and network latencies. Encrypting the link over which credentials are transmitted is a common method used to prevent cracking material from falling into the wrong hands. This could be done using a virtual private network (VPN) or secure sockets layer (SSL). In this way, a system may continue to use password systems which are vulnerable to brute-force attacks but trust the network to protect against them.

Kahoot! ASA et al. Ex. 1024-013

4.2 Hardening Password Systems

In many cases, protecting password material can be a complicated or impossible task thus, the security of the system lies in the difficulty of offline brute-force attacks. A number of methods for making password guessing more difficult have been proposed. One proposed solution is simply to make the encryption more complex and computationally expensive, thus reducing the speed of brute-force attacks. However, this presents a few problems. First, this may put a great strain on the server. In the case of Kerberos, if the KDC must hash TGTs 1000 times instead of just once, the computational load on the server would increase considerably. This exposes authentication servers to DoS attacks, since an attacker can repeatedly attempt to authenticate in most systems. Complex encryption would also make it difficult to incorporate legacy systems into new authentication schemes. For example, an older 100MHz Pentium system cannot do 100 billion MD5 hashes very easily. Moreover, low power devices are also much less capable of complex encryption. Other solutions, such as hardware-dependent encryption algorithms [28], result in a hardware-cryptography arms race. As hardware increases in speed, cryptography is deliberately slowed to maintain its security.

Another often proposed solution is to increase the minimum number of characters required for passwords. Unfortunately, this solution is restricted by a fundamental human limitation of remembering no more than 7 random items easily. While some users may be able to memorize random strings of 12 or more characters, many will not be able to, and will be forced to write down passwords or pick passwords will very little randomness. A good password recovery tool will be able to try such non-random combinations quickly, negating most of the benefit to having a longer password.

Implementation restrictions also make this solution practically difficult. Any system that must inter-operate with older crypt() implementations is limited to 8 characters. Many sites today require passwords of no longer than 8 characters because of this. Additionally, since users must remember such a large number of passwords, they often re-use them from site to site. As such, they often pick passwords that will be universally accepted, thus restricting their password choices to 6-8 character passwords conforming to standard password policies.

Pass phrases are an interesting alternative to passwords. In this system, a user would choose a pass phrase of around 7 words. Since there are around 500,000 words in the English language [4], the potential combination of these words can provide a larger password space. However, in practice, informal studies have shown that it may be the case that pass phrases often provide less randomness than passwords [29]. This is conjectured to be the result of users picking common phrases and words, resulting in less total randomness.

Two-factor authentication is another solution that has gained recent prominence with the strong recommendation of the Federal Financial Institutions Examination Council that two-factor authentication should be used by 2006 for all Internet financial transactions [15]. As long as the method of two-factor authentication used includes some type of random number or symmetric key, this information could be combined with the user's password to create a key that falls randomly within the keyspace of the encryption method used. This then eliminates the fundamental restriction of passwords: that they only occupy a very limited subset of the keyspace supported by the encryption used.

Kahoot! ASA et al. Ex. 1024-014

5 Related Work

Morris and Thompson first addressed the issue of password security in 1979 [26] by describing the challenges faced by the UNIX password system. They observed problems that existed within the system stemming from the availability of the password file and then identified guessing passwords as a general approach that was successful in penetrating the system. However, the time to encrypt each guess and compare the result with the file entries was highlighted as the main challenge in password guessing. They analyzed passwords from 1 to 6 characters long from key spaces of 26 to 128 characters and found that exhaustively searching the key space was beneficial in finding a fraction of a system's passwords given enough time. To simplify the searching task, they also noticed that the users of the system chose short and simple passwords, which greatly reduced the key space.

In order to make cracking user passwords more challenging, Morris and Thompson proposed a list of tips to make stronger passwords. The suggestions were attempts to slow the process of password cracking and included basic ideas like choosing longer passwords, choosing passwords constructed from a larger character set, or having the system create passwords. The authors also proposed password salting, combining the password with extra well-known data, as a technique to make pre-computation impossible and increase the time necessary to crack a password. These defenses became the basis for future password cracking prevention techniques.

Ten years later, a paper was published discussing the claims of Morris and Thompson as well as the progress of password security and cracking [19]. Like its predecessor, the paper examined the performance of key space searches. They looked at the possible times for cracking passwords with the same key space as Morris and Thompson, but examined lengths ranging from 4 to 8 characters. With the addition of password salting, the searches had indeed become more complex. The authors claimed that a large key space of 95 characters "is large enough to resist brute force attacks in software ... It is impossible to use exhaustive search over the large search space..." However, they determined that password cracking was very possible if the search space was limited. This could be done by creating a common password word list to guess passwords from instead of attempting to guess every possible password.

In order to maximize the difficulty of password cracking, [19] discussed execution speed of the hashing mechanism and password entropy. The authors concluded that, because computing speed and power were changing, attacking the problem by increasing the speed of the encryption algorithm was not plausible. They also analyzed other solutions including changing the encryption algorithm and better protecting the cracking material. It was concluded that making passwords less predictable was the principal defense against password cracking.

Dictionary attacks are the fastest, easiest way to crack passwords because passwords are commonly chosen from a small set of words. In order to prevent these fast, simple attacks, systems implemented policies that required passwords contain a certain amount of entropy. The policies include rules on minimal length and required password characters. To enforce these policies, password checking software was developed, which determined if a given password had enough entropy to be considered secure. However, dictionary attacks then evolved by exploiting common non-dictionary choices for

Kahoot! ASA et al. Ex. 1024-015

passwords. The techniques used by these attacks included searching for random capitalization, permutations of dictionary words and usernames, letter and number manipulations, and foreign language words. These attacks continue to evolve by examining and exploiting common policies. Unfortunately, research has shown that despite password policy advice, users still tend towards dictionary words for passwords [22].

Sophisticated analysis of the English language has aided in password guessing. For example, character frequency, once very successfully used as a spellchecker in UNIX, is now being used in password cracking [3]. Analysis of common passwords has also contributed to faster password cracking. Possible passwords are tried in a certain order based on how common the password is. From these advanced methods, we see that password guessing techniques continue to evolve as long as passwords are still in use. As a result, a variety of solutions have been proposed to combat password guessing.

Twenty five years after Morris and Thompson's paper, modern passwords are still vulnerable to offline cracking attacks. Basic hashes and digests are still used to encrypt these passwords, thus today's cracking material is similar to that available in the 1980s. However, the ability hash passwords, and thus recover passwords, has drastically improved due to developments in software that have quickened the performance of these encryption techniques, sometimes by as much as a factor of 5 [13]. These improvements have impacted the speed at which passwords can be cracked, thus increasing the difficulty of preventing offline password cracking.

6 Conclusion

The limited protection passwords provide to authentication systems that allow offline attacks is clearly no longer sufficient to resist serious attacks. Such systems are fundamentally restricted in the amount of protection they can provide, while the resources of the attackers grows exponentially. Due to the ease with which even random passwords will be recoverable in the next 5 years, the security of any system based on passwords will be equivalent to the availability of the cracking material, not how random the passwords are. As such, protocols must be designed to not allow any type of offline attack, and the material that can be used to mount such an attack must be protected with the understanding that its confidentiality is equivalent to the security of the authentication mechanism as a whole.

References

1. AMD 3-year technology outlook. http://www.amdcompare.com/techoutlook/.
2. EFF DES cracker project. www.eff.org/Privacy/Crypto/Crypto_misc/DESCracker/.
3. John the ripper password cracker. http://www.openwall.com/john/.
4. Number of words in the English language. http://hypertextbook.com/facts/2001/JohnnyLing.shtml.
5. pwgen CVS changelog. http://pwgen.cvs.sourceforge.net/pwgen/src/ChangeLog?revision=1.8&view=markup.
6. pwgen password generator. http://sourceforge.net/projects/pwgen/.
7. Sans institute. http://www.sans.org/aboutsans.php.

Kahoot! ASA et al. Ex. 1024-016

8. SANS password policies. http://www.sans.org/resources/policies/.
9. Unix cyrpt man page. http://bama.ua.edu/cgi-bin/man-cgi?crypt_unix+5.
10. The magical number seven, plus or minus two: Some limits on our capacity for processing information, 1956.
11. Mihir Bellare, David Pointcheval, and Phillip Rogaway. Authenticated key exchange secure against dictionary attacks. *Lecture Notes in Computer Science*, 1807:139–156, 2000.
12. S. M. Bellovin and M. Merritt. Limitations of the kerberos authentication system. *SIGCOMM Comput. Commun. Rev.*, 20(5):119–132, 1990.
13. Eli Biham. A fast new DES implementation in software. *Lecture Notes in Computer Science*, 1267:260–272, 1997.
14. Shekhar Y Borkar, Pardeep Dubey, Kevin C Kahn, David J Kuck, Hans Mulder, Stephen S Pawlowski, Justing R Rattner, R M Ramanathan, and Vince Thomas. Platform 2015: Intel Processor and Platform Evolution for the Next Decade. Technical report, Intel, 2005.
15. Federal Fiancial Institutions Examination Council. Authentication in an internet banking environment. http://federalreserve.gov/boarddocs/srletters/2005/SR0519a1.pdf.
16. Larry Cuban. *Oversold and Underused: Computers in the Classroom*. Harvard University Press, 2001.
17. Manek Dubash. Moore's Law is dead, says Gordon Moore. *http://www.techworld.com/opsys/news/index.cfm?NewsID=3477*, 2005.
18. Magnus Ekman, Fredrik Warg, and Jim Nilsson. An in-depth look at computer performance growth. *SIGARCH Comput. Archit. News*, 33(1):144–147, 2005.
19. David C. Feldmeier and Philip R. Karn. Unix password security - ten years later. In *CRYPTO '89: Proceedings of the 9th Annual International Cryptology Conference on Advances in Cryptology*, pages 44–63, London, UK, 1990. Springer-Verlag.
20. Radhakrishna Hiremane. From Moore's Law to Intel Innovation - Prediction to Reality. *Technology@Intel Magazine*, April 2005.
21. Ian Jermyn, Alain Mayer, Fabian Monrose, Michael Reiter, and Aviel Rubin. The Design and Analysis of Graphic Passwords. In *Proceedings of the 8th Annual USENIX Security Symposium*, 1999.
22. Daniel V. Klein. "foiling the cracker" – A survey of, and improvements to, password security. In *Proceedings of the second USENIX Workshop on Security*, pages 5–14, Summer 1990.
23. J. Kohl and C. Neuman. RFC 1510: The Kerberos Network Authentication Service (V5), September 1993. Status: PROPOSED STANDARD.
24. Rob Lemos. Passwords: The Weakest Link. http://news.com.com/2009-1001-916719.html, 2002.
25. Fabian Monrose and Aviel Rubin. Authentication via Keystroke Dynamics. In *Proceedings of the 4th ACM Conference on Computer and Communication Security*, 1997.
26. Robert Morris and Ken Thompson. Password security: a case history. *Commun. ACM*, 22(11):594–597, 1979.
27. Arvind Narayanan and Vitaly Shmatikov. Fast dictionary attacks on passwords using time-space tradeoff. In *CCS '05: Proceedings of the 12th ACM conference on Computer and communications security*, pages 364–372, New York, NY, USA, 2005. ACM Press.
28. Niels Provos and David Mazières. A Future-Adaptable Password Scheme. In *USENIX Annual Technical Conference, FREENIX Track*, pages 81–91, 1999.
29. Arnold G. Reinhold. Results of a survey on pgp pass phrase usage. http://www.ecst.csuchico.edu/ atman/Crypto/misc/pgp-passphrase-survey.html.
30. R. Rivest. The MD5 Message-Digest Algorithm . RFC 1321 (Informational), April 1992.
31. Wayne C. Summers and Edward Bosworth. Password policy: the good, the bad, and the ugly. In *WISICT '04: Proceedings of the winter international synposium on Information and communication technologies*, pages 1–6. Trinity College Dublin, 2004.

Kahoot! ASA et al. Ex. 1024-017

32. W. Timothy Polk William E. Burr, Donna F. Dodson. Electronic authentication guidelines. NIST Special Publication 800-63.
33. Thomas Wu. The secure remote password protocol. In *Proceedings of the 1998 Internet Society Network and Distributed System Security Symposium*, pages 97–111, 1998.
34. Thomas Wu. A real-world analysis of Kerberos password security. In *Internet Society Network and Distributed System Security Symposium*, 1999.

Appendix

Unix Crypt

The previous work in this field has examined password cracking primarily as it applies to cracking attempts against the Unix/Linux /etc/passwd, so we start by examining this password storage type to give a sense of how modern techniques and equipment compare to what was previously available. This type of authentication system is used to authenticate users to a Unix/Linux system. In the traditional Unix crypt system, hashes of users' passwords are stored in a password file often named either /etc/password or /etc/shadow, with the 2 letter salt prepended to the hash. A user enters their password, which is then combined with the salt in the password file, and then encrypted using a variation of the DES algorithm. The resulting ciphertext is compared with the hash in the password file, and if the values match, the user is successfully authenticated. Newer systems, such as the one first found in modern crypt function, hash the password with MD5 repeatedly (up to 1000+ times), instead of just once[28].

For this type of authentication system, an attacker must somehow obtain a copy of this password hash file. Unfortunately, this can be made available to an attacker in a variety of ways. The simplest of these is if an attacker has root access on a machine, in which case he can simply copy the /etc/shadow file. If the password hashes have not been moved to /etc/shadow, they will reside in the world-readable /etc/passwd file, in which case an attacker with normal user access to the system can simply copy the file. However, amongst other attacks, there is one attack in particular which allows a large number of attackers access to the password file. The Network Information Services, or NIS, is often used to centralize authentication decisions over a large number of machines. NIS provides a utility, ypcat, which allows users to view portions of information about system users. We found ypcat to be often misconfigured in a way that allows any user on any system connected to NIS to simply ypcat the password hash portion of each user in the system. In this way, an attacker can gain access to the credentials of each user on any system tied to NIS.

The actual process of guessing a user's password is very simple. To recover passwords from this password file, an attacker takes candidate passwords, combines them with the appropriate salt, which is well known, and applies the appropriate hashing technique to this value. The attack then checks to see if the result from hashing his guess matches the hash value in the password file.

Kerberos

We also evaluate password cracking as it relates to modern versions of Kerberos. Kerberos, a popular single-sign-on protocol, is widely touted today as a solution to "the

Kahoot! ASA et al. Ex. 1024-018

password problem." It is used to authenticate to a variety of services, including IPsec, Email, Web Services, Directory Services, and many more. Because Kerberos is often used as a single-sign-on service, a compromise of Kerberos credentials is often equivalent to a compromise of the users' credentials to every service in the network. Unfortunately, Kerberos, in every version, is vulnerable to a variety of password-guessing attacks[12,34].

One of the biggest issues with Kerberos as it relates to password cracking is that as opposed to most Unix/Linux systems, where an attacker must have a valid user account (or have compromised one), all the cracking material necessary to mount an offline attack against Kerberos credentials can be obtained either by anyone who asks or anyone who can sniff Kerberos traffic, depending on the restrictions in place. During a client's initial authentication In the Kerberos protocol, a client sends an authentication request to a server in charge of authentication for the Kerberos realm called the KDC. If the client makes a correct request, the KDC will return a token called a ticket granting ticket (TGT). This token can be used to obtain credentials to any Kerberized service the client can access. Unfortunately, when this TGT is given to the client, it is transmitted over the network, encrypted with a key derived from the user's password. While the user's password itself is never sent in any form, this TGT is still vulnerable to password guessing attacks, as described below.

An attacker has a variety of options for obtaining cracking material (the TGT) required for this attack. In Kerberos v4, a KDC will return a TGT to anyone who asks for it. Thus, in this case, an attacker's job is completely trivial, and he can easily obtain a TGT to crack for each user in the system by simply asking. However, Kerberos v5 introduced the idea of preauthentication. With preauthentication, a user must use the key derived from his password (as described above) to encrypt a timestamp, which is included in the client's request for TGT. The server will only return a TGT if the timestamp received by the server decrypts correctly with the client's key. In this way, the server attempts to insure that a TGT is only sent to the user to whom it belongs.

However, an attacker attempting to crack a Kerberos 5 deployment still has a number of options for recovering a TGT. First, many Kerberos deployments do not have preauthentication required for all users. In this situation, an attacker may simply ask for TGTs as he did for Kerberos v4. Many deployments, in order to ensure backwards compatibility, still support Kerberos v4, so an attacker may simply ask for v4 tickets for each user. Finally, in any of these systems, the TGT itself is sent over the network in the clear, so an attacker that can sniff the network can trivially recover the TGT.

In order to compromise Kerberos credentials, an attacker first captures the TGT using one of the aforementioned methods. Then, an attacker generates a password guess. This guess is transformed into a key using the Kerberos "stringToKey" function, which uses both the password guess and information found in the TGT itself, such as the user's name and the name of the Kerberos realm. Then, this key is used to decrypt the captured TGT. Since each TGT, if decrypted correctly, contains the string "krbtgt", it is easy for an attacker to know if the decryption, and therefore the candidate password, was correct.

Kahoot! ASA et al. Ex. 1024-019