TwoKind Authentication: Protecting Private Information in Untrustworthy Environments

Katelin Bailey, Apu Kapadia, Linden Vongsathorn, Sean W. Smith

Department of Computer Science
Dartmouth College
Hanover, NH 03755, USA

{katelin, akapadia, linden, sws}@cs.dartmouth.edu

ABSTRACT

Users often log in to Internet sites from insecure computers and more recently have started divulging their email passwords to social-networking sites, thereby putting their private communications at risk. We propose and evaluate *TwoKind Authentication*, a simple and effective technique for limiting access to private information in untrustworthy environments. In its simplest form, *TwoKind* offers two modes of authentication by providing a `low` and a `high` authenticator. By using a `low` authenticator, users can signal to the server that they are in an untrusted environment, following which the server restricts the user's actions.

We seek to evaluate the effectiveness of multiple authenticators in promoting safer behavior in users. We demonstrate the effectiveness of this approach through a user experiment — we find that users make a distinction between the two authenticators and generally behave in a security-conscious way, protecting their `high` authenticator the majority of the time. Our study suggests that *TwoKind* will be beneficial to several Internet applications, particularly if the privileges associated with the `low` authenticator can be customized to a user's security preferences.

Categories and Subject Descriptors

K.6.5 [**Management of Computing and Information Systems**]: Security and Protection—*Authentication*; H.1.2 [**Information Systems**]: User/Machine Systems—*Human factors*

General Terms

Security, Human Factors

Keywords

authenticators, passwords, principle of least privilege

WPES'08, October 27, 2008, Alexandria, Virginia, USA.
Copyright 2008 ACM 978-1-60558-289-4 /08/10 ...$5.00.

1. INTRODUCTION

Users of online applications are increasingly placing their private information at risk; a large number of users routinely access Internet sites from untrustworthy computers such as email kiosks and Internet cafes. Malicious administrators of such computers, and even users who are able to install rogue applications, can easily compromise a user's session and gain unauthorized access to private data. A user's session can be hijacked, or kept alive by spoofing the logout screen, following which the attacker has unfettered access to all the user's private emails, personal profile information, and so on. In cases where passwords are used as authenticators, the potential damage is even worse because the user's authenticator can be compromised and saved for later use.

The risk of session or authenticator compromise is not limited to the use of untrustworthy computers. Social-networking services such as Facebook[1] ask for a user's login information for external email services such as Google Mail.[2] These social-networking services proceed to download the user's address book and use it to find the user's existing contacts in the social network. In these situations, the user provides an online application with unnecessarily unrestricted access to all of the user's private data on another application. Ideally, the user would authorize the service to download only the user's address book, and disallow access to email. To address these issues, we propose *TwoKind Authentication*, an authentication technique that allows users to *limit the capabilities of an authenticated session*, thereby limiting the amount of damage that can be caused by session or authenticator compromise.

Current authentication mechanisms such as one-time passwords [8] [11] (e.g., RSA SecurID), privileged "trading passwords" (such as those used by eTrade [5] while placing trades), or even PKI tokens do not fully solve this problem. One-time passwords limit the future damage possibly caused by stolen credentials but allow full-scale compromise in a single hijacked session. PKI tokens do not protect against hijacked sessions either, and they can be susceptible to authenticator hijacking [10]. eTrade-style trading passwords are required by server policy, where users must re-enter a trading password while executing privileged actions such as trades. Such systems, however, have usability concerns, since the default mode of access is that of low privilege. Re-

[1] http://www.facebook.com

[2] http://mail.google.com

quiring users of an email application to enter a high-privilege password each time they want to access archived email, for example, would be a nuisance. *TwoKind* does not prevent a session from being hijacked; rather, it gives users a convenient method to effectively limit the damage caused by hijacking, and allows more usable access modes from trustworthy environments.

In its simplest form, *TwoKind* features two modes of authentication — `high` and `low`. *TwoKind* authenticators could include passwords, PKI-based keys, or hardware tokens. To signal untrustworthy environments to the server, users employ their `low` authenticator to limit the privileges of the session. For example, a user's `low` authenticator for an email service may allow access to only the user's new messages (and not previously viewed messages, messages in folders, and so on); the `low` authenticator for a bank account may disallow any financial transactions or access to financial records other than the account balance. *TwoKind* thus allows users to log in with full-privileged access under normal circumstances, making it less intrusive in general. More generally, *TwoKind* allows users to assign specific permissions to their `low` authenticator, or to use any number of `low` authenticators with different, but limited, capabilities.

We evaluate *TwoKind*'s effectiveness in a general population and seek to determine how often users will protect their `high` authenticator in unsafe environments. We focus on the question of whether users would apply such a bimodal cost-benefit tradeoff for authenticated sessions.[3] Our user experiment showed that 70% of subjects were able to apply *TwoKind* effectively, i.e., they made pragmatic use of *TwoKind* based on their assessment of risk. Additionally, 49% of the time, subjects were able to protect their `high` passwords in unsafe environments, and we therefore propose *TwoKind* as a useful authentication method.

Paper outline.

The remainder of the paper is structured as follows. We describe *TwoKind Authentication* in Section 2. Section 3 describes our methodology for evaluating *TwoKind Authentication*, and in Section 4 we detail the results of our experiment. We present related work in Section 5, future work in Section 6, and conclude in Section 7. An extended version of this paper, with more details on the user study and its results, is available as a Dartmouth Computer Science Technical Report [1]. We note that our proposed user experiment was presented as a poster at SOUPS 2007 [2].

2. TWOKIND AUTHENTICATION

We now describe *TwoKind* more precisely, along with its generalization to more than two authenticators.

2.1 Two authenticators for the same account

In *TwoKind*, users are assigned two authenticators, `high` and `low`, for the *same account*,[4] where the `low` authenticator is associated with restricted privileges. In PKI-based *TwoKind*, a user would carry two PKI tokens, and use the `low` token in unsafe environments, and in password-based *TwoKind*, users would use a `low` password. These authenticators are used to signal to the server whether the user is in a *safe* or *unsafe* situation, depending on whether they trust the security of the session. For example, a user may determine that using an email kiosk is unsafe, or that giving Facebook his or her Google Mail password is unsafe.

Let A be the set of all privileges for user u, and $P(x)$ be all the privileges associated with the authenticator x. In our model, we assume that the `high` authenticator is the default authenticator as would be used without multiple authenticators, and has associated with it all the privileges $P(\texttt{high}) = A$. The `low` authenticator has some proper subset of these privileges, i.e., $P(\texttt{low}) \subset A$. The privileges associated with the session are determined by the `low` or `high` authenticator that is used, resulting in privileges $P(\texttt{low})$ or $P(\texttt{high})$ respectively. We note that either the server or the user can define the set of privileges $P(\texttt{low})$ associated with the `low` authenticator, although we do not examine the usability of user-defined privileges in this paper.

2.2 Multiple authenticators

Although we focus on and evaluate the simpler concept of *TwoKind Authentication*, we expect that some users may desire the ability to create several authenticators for different uses. For example, the user may create an "address-book password" so that social-networking sites may access the user's address book, but nothing else connected with that email account. The user may create a separate "travel password" for use in Internet cafes, allowing the sending and receiving of new mail only. For n-Kind authentication, users possess the following authenticators: $\texttt{low}_1, \texttt{low}_2, \ldots, \texttt{low}_{n-1}, \texttt{high}$. For each $\texttt{low}_i$, we have that $P(\texttt{low}_i) \subset A$, and these sets are not necessarily disjoint (different `low` passwords may have some privileges in common). Again, either the server or the user could define these sets of privileges. The server may even allow a combination of the two (a static set of server-defined permissions, with the option of modifying these permissions).

3. USER EXPERIMENT

We designed a user study to determine how users would employ their *TwoKind* authenticators when presented with safe and unsafe environments. Subjects participated in a game designed to measure their risk-taking habits when given *TwoKind* as a means to mitigate the risks.

3.1 Experimental protocol

Because this study was being performed on college students, we designed a game emulating a Facebook-style application. Subjects were told they were in one of two environments and were presented with a series of "desired updates" to their profile. To access their information in either of these environments, subjects were given the option to log in to an environment by indicating use of a `high` password, a `low` password, or to not log in at all. When logged in, the subject was presented with a screen of their information, with certain editing capabilities removed if they were logged in with their `low` password. Figure 1 shows a screenshot of *Green Book Online*. For an in-depth discussion of the study's design, see the extended version of this paper [1].

3.2 Task completion and points

There were two types of updates a subject could perform: updates that could be completed with either password, and

[3] We do not investigate the usability of a particular *instantiation* such as password-based *TwoKind*, where the usability of passwords would interfere with our measurements.

[4] As opposed to two *root* and *user* accounts.

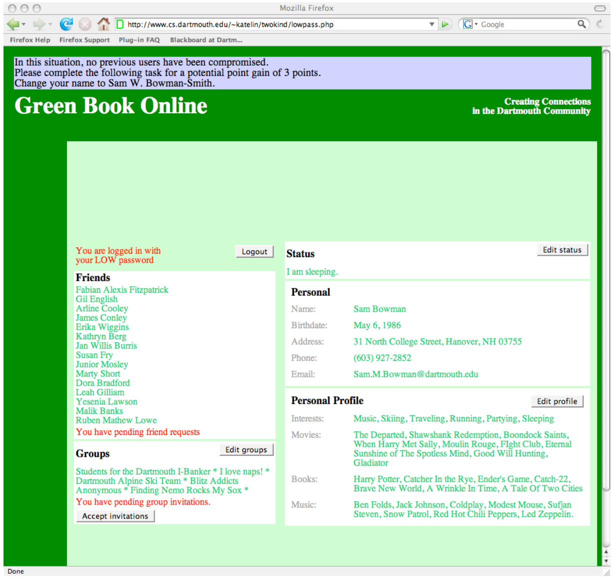

In this situation, no previous users have been compromised.
Please complete the following task for a potential point gain of 3 points.
Change your name to Sam W. Bowman-Smith.

Green Book Online

Creating Connections
in the Dartmouth Community

You are logged in with
your LOW password

Friends

Fabian Alexis Fitzpatrick
Gil English
Arline Cooley
James Conley
Erika Wiggins
Kathryn Berg
Jan Willis Burris
Susan Fry
Junior Mosley
Marty Short
Dora Bradford
Leah Gilliam
Yesenia Lawson
Malik Banks
Ruben Mathew Lowe

You have pending friend requests

Groups

Students for the Dartmouth I-Banker * I love naps! * Dartmouth Alpine Ski Team * Blitz Addicts Anonymous * Finding Nemo Rocks My Sox *

You have pending group invitations.

Status

I am sleeping.

Personal

Name:	Sam Bowman
Birthdate:	May 6, 1986
Address:	31 North College Street, Hanover, NH 03755
Phone:	(603) 927-2852
Email:	Sam.M.Bowman@dartmouth.edu

Personal Profile

Interests:	Music, Skiing, Traveling, Running, Partying, Sleeping
Movies:	The Departed, Shawshank Redemption, Boondock Saints, When Harry Met Sally, Moulin Rouge, FIght Club, Eternal Sunshine of The Spotless Mind, Good Will Hunting, Gladiator
Books:	Harry Potter, Catcher In the Rye, Ender's Game, Catch-22, Brave New World, A Wrinkle In Time, A Tale Of Two Cities
Music:	Ben Folds, Jack Johnson, Coldplay, Modest Mouse, Sufjan Steven, Snow Patrol, Red Hot Chili Peppers, Led Zeppelin.

Figure 1: Screenshot of the `low` privilege environment in *Green Book Online*, where buttons for modifying personal information and friends have been removed so that it is absolutely clear that users may not access these functions.

updates that required their `high` password. The subject gained 3 points for completing a high-privilege task and 1 point for completing a low-privilege task. There was an unspecified probability that if a subject logged in to an unsafe environment with either password, they would lose some of the points they had accumulated. A compromised `high` password lost 6 points, while a compromised `low` password lost 2 points. At the end of the game, the subject was given a certain amount of money directly related to their final score.

For a fixed probability p of compromise, the expected gain in points for unsafe environments is $1-2p$ and $3-6p$ for low- and high-privilege tasks respectively. The value of p is unspecified to avoid biasing users towards either always risking ($p \leq 1/2$) or never risking ($p > 1/2$) their authenticators. Furthermore, we did not want to influence users to employ one authenticator over another. Thus for a fixed p, there is either an expected gain or expected loss for both types of tasks. The variance, however, differs for the two situations, and thus we can study how many users tend to take higher risks for potentially higher rewards. This game, therefore, allowed us to study the effectiveness of the *TwoKind* model without obviously biasing the subjects' choices.

3.3 Survey

To further explore participants' reactions to the TwoKind method, and to collect information such as whether participants had a background in computer science, we administered a closing survey to each participant.

4. RESULTS

We now describe the results of our study. We identify several categories of users depending on their behavior in various situations and present results of our survey questions.

4.1 General Patterns of Behavior

While designing our study, we expected subjects to react to both the type of environment (safe or unsafe) and the level of privilege (`high` or `low`) required for each desired update. We found that 79% of all subjects fit this pattern. Please refer to the extended version of this paper [1] for the specific types of behaviors; here we will discuss the overarching trends that subjects followed.

In the study, all but four users followed the principle of least privilege, using their `high` password only when it was necessary to complete the task, and using the `low` password whenever they could. The prominence of this behavior is a positive sign towards individuals using the lowest privilege possible, and by proxy leaving themselves open to as little risk as they can. It is interesting to see how these users behave in unsafe environments, since the principle of least privilege by itself will result in the compromise of the `high` authenticator when used in unsafe environments.

While few users were willing to skip tasks on a consistent basis, they were able to show good decision-making skills in weighing the risks of using their `high` password as opposed to the `low` password, often protecting their `high` password more than their `low` password. This behavior is consistent with the goals of *TwoKind*, which aims not to create a method of absolute rules for when to use `high` versus `low`, but to give the users the choice of protecting some capabilities over others, and allowing low-privilege tasks to be accomplished in unsafe environments.

Overall, 70% of subjects were sensitive to both environment and privilege and made a distinction between the passwords in an *unsafe* environment (risking `high` less). We conclude that *TwoKind* provides a multiple authentication method that is useful to this 70% of users, who make a conscious decision to differentiate between passwords based on the perceived risks and are more protective of their `high` authenticator.

4.2 Task Groups

It is also interesting to see how subjects reacted to groups of tasks. There are four main groups that tasks fall into: permutations of the safe or unsafe environment combined with high- or low-privilege tasks. Figure 2 shows how subjects reacted to these four situations.

1. *High-privilege task, safe environment.* Subjects chose to log in with their `high` password an overwhelming majority (90%) of the time, since there was no reason to do otherwise.

2. *High-privilege task, unsafe environment.* We observed a fairly even split between users who attempted to log in with their `high` password and those who skipped the task, with a small percentage trying multiple authenticators. These users were probably trying to accomplish the task with the `low` password, and most of them skipped when this was not possible. About half (49%) of the time subjects would have protected their `high` password from compromise.

3. *Low-privilege task, safe environment.* Although in this environment it would have been safe to use the `high` password, we see subjects choosing to use their `low` password an overwhelming majority (86%) of the

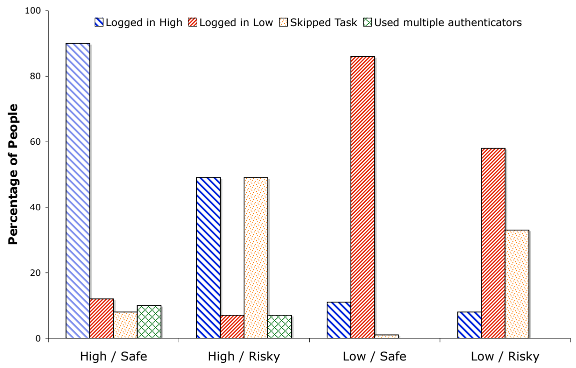

Figure 2: Chart of task groupings. This figure shows the overall reaction of subjects to various permutations of safe/risky environment and high/low-privilege tasks.

time, which shows that subjects followed the principle of least privilege.

4. *Low-privilege task, unsafe environment.* A small percentage (8%) of the time subjects risked their `high` password in this environment, but the majority of tasks (58%) were accomplished with their `low` password. It is interesting to see that a significant amount of the time (33%) subjects chose to also protect their `low` password from compromise as well, and skipped these tasks.

It is important to remember that each of these trend groups draws conclusions from 165 separate instances of task completion: five tasks each completed by 33 subjects, which means that individual users could have been inconsistent over these groups. In Section 4.1, we placed subjects into patterns that varied widely between users. In Section 4.2 the overall behavioral trends give us an idea of how the method would be used by a larger population. Together, these patterns from specific users and behavioral trends over groups of tasks give a full picture of users' behavior. The trends show that 49% of the time subjects protect their `high` password from compromise, and the patterns show that a higher percentage of users (70%) meet the goals of *TwoKind*: making pragmatic decisions depending on the type of environment. We posit that this is a large enough portion of the population to justify *TwoKind* as a viable future authentication mechanism.

4.3 Survey Results

We administered a closing survey to the subjects, asking them a few questions, which have helped us to identify where *TwoKind* may be most useful. Please refer to the extended version of this paper [1] for a complete description of the survey. Table 1 shows the responses received from two questions that asked whether the subjects would use *TwoKind* for BlitzMail, an email application that is pervasive on the Dartmouth College campus, or for "friend finder" applications offered by websites such as Facebook.

In general, subjects were more receptive to the idea of using *TwoKind* for "friend finder" applications. Many subjects expressed concern about giving their passwords to these applications, and the majority of those who do provide their login information were interested in having the ability to provide a `low` password instead of their main password. Users were also receptive to the idea of using *TwoKind* for email, but desired more flexibility in the permissions assigned to each authenticator than was suggested in the survey. Thus, we conclude that a more advanced form of *TwoKind*, with flexible permissions, would be appealing to the general population.

5. RELATED WORK

TwoKind Authentication differs from the principle of least privilege [12] in that we expect users to employ authenticators based on the trustworthiness of the environment. For example, users may trust their personal laptop machines, and log into their email account with full-privileged access, but use a `low` password when using an untrusted machine.

Instead of relying on the user to signal the type of environment, the server may be able to determine the trustworthiness of the client through remote attestation [13, 4, 6]. The problem with these techniques is that they have not yet seen widespread deployment. Furthermore, even if the server has determined that the remote platform is untrustworthy, it has no way of warning the user, thereby limiting the utility of these approaches. In contrast, TwoKind allows users to determine the trustworthiness of the environment for themselves and act upon that determination.

Other work has explored "proxy certificates" for delegating limited privileges to other users or platforms [3, 9, 14]. While these techniques are certainly viable options for untrustworthy environments, they require modifications to the client machines. Again, TwoKind does not require any such modifications, and is a practical solution in the short term.

One concern with our approach is that it requires users to memorize additional passwords. On the issue of memorability, Yan et al. [15] have found that passwords based on mnemonic phrases tend to be as memorable as naively-selected passwords, and as secure as randomly-selected passwords. Gaw and Felten [7] found that users tend to have only about three distinct passwords that they reuse across accounts. We expect that users may create a `low` password that is reused across several accounts, and will have a low-level of security associated with that password (further justifying its reuse across several accounts).

6. FUTURE WORK

We performed our study on 33 Dartmouth College undergraduate students, and our results are therefore representative of a college-aged demographic. Further work could explore the general public's approach to *TwoKind* and how it differs from our sample.

The results from our study suggest several additional areas of further investigation. It would be interesting to study whether having multiple passwords decreases each individual password's security, whether passwords chosen for a particular account are related (does a user's `low` password provide a hint about the `high` password?), and whether users would react to having to remember and maintain additional passwords. Our study addresses short actions, where subjects only have to complete one task. Investigating behavior when users log into their accounts for longer sessions may

Table 1: Responses to the survey questions

Question	*No*	*Yes*	*Uncertain*
Would use TwoKind *with BlitzMail*	55%	36%	9%
Would use TwoKind *with "friend finder"*	45%	48%	6%

yield different results. Further exploration of the usability of other authentication methods like PKI tokens would provide a more complete survey of security solutions.

Additionally, it would be interesting to study how privilege levels could be set and maintained, and the behavior of users when given the opportunity to set their own privilege levels. In our study, we told users whether a situation was safe or unsafe. Research into users' ability to judge the security of real-world situations would provide insight into the effectiveness of solutions like TwoKind.

7. CONCLUSIONS

We proposed a method called *TwoKind Authentication*, which protects users from malicious administrators or thirdour party services. Using a `low` authenticator, users can signal untrustworthy environments to the server and reduce the privileges associated with that session. A compromised authenticator, therefore, allows attackers only limited access to private information. We performed a user experiment with 33 subjects, in which 70% of users employed the two authenticators in a way that was consistent with the goals of *TwoKind*, including making distinctions between environments, recognizing privilege levels, and protecting the `high` authenticator by means of the `low` authenticator. Furthermore, 49% of the time, subjects did not risk their `high` authenticator in unsafe environments.

We believe that *TwoKind* is a useful authentication method, which is an improvement on the current practice of either using a single high-privilege authenticator, or repeatedly requiring high-privilege authenticators for certain actions. Although users may not always use the *TwoKind* method ideally, allowing their `high` authenticator to be compromised on occasion, it appears that the majority of users would employ *TwoKind* to their benefit. Since the study on the whole demonstrated that users are willing to use the `low` authenticator to protect the `high` authenticator, *TwoKind* seems to generally increase the security of high-privilege actions and reduce the risk of compromise in unsafe environments.

8. ACKNOWLEDGMENTS

We thank Denise Anthony, Sergey Bratus, Chris Masone, Sara Sinclair, and the anonymous reviewers and pre-testers of our study for their helpful comments. This research was supported in part by the NSF, under Grant CNS-0448499, the Bureau of Justice Assistance, under grant 2005-DD-BX-1091, and the Women in Science Project at Dartmouth College. The views and conclusions do not necessarily reflect the views of the sponsors.

9. REFERENCES

[1] K. Bailey, A. Kapadia, L. Vongsathorn, and S. W. Smith. TwoKind authentication: Protecting private information in untrustworthy environments (extended version). Technical Report TR2008-632, Dartmouth College, Aug 2008.

[2] K. Bailey, L. Vongsathorn, A. Kapadia, C. Masone, and S. W. Smith. TwoKind authentication: Usable authenticators for untrustworthy environments (poster abstract). In *Symposium on Usable Privacy and Security (SOUPS 2007)*, pages 169–170, July 2007.

[3] J. Basney, M. Humphrey, and V. Welch. The MyProxy online credential repository: Research articles. *Softw. Pract. Exper.*, 35(9):801–816, 2005.

[4] E. Brickell, J. Camenisch, and L. Chen. Direct anonymous attestation. In *CCS '04: Proceedings of the 11th ACM conference on Computer and communications security*, pages 132–145, New York, NY, USA, 2004. ACM.

[5] eTrade Trading Passwords. https://www.etradeaustralia.com.au/EStation/hep_aec_connecting.asp.

[6] S. Garriss, R. Caceres, S. Berger, R. Sailer, L. van Doorn, and X. Zhang. Towards trustworthy kiosk computing. *Mobile Computing Systems and Applications, 2007. HotMobile 2007. Eighth IEEE Workshop on*, pages 41–45, 8-9 March 2007.

[7] S. Gaw and E. W. Felten. Password management strategies for online accounts. In *SOUPS '06: Proceedings of the second symposium on Usable privacy and security*, pages 44–55, New York, NY, USA, 2006. ACM.

[8] L. Lamport. Password authentication with insecure communication. *Communications of the ACM*, 24(11):770–772, Nov. 1981.

[9] J. Marchesini and S. W. Smith. SHEMP: Secure Hardware Enhanced MyProxy. In *PST*, 2005.

[10] J. Marchesini, S. W. Smith, and M. Zhao. Keyjacking: the surprising insecurity of client-side SSL. *Computers and Security*, 24(2):109–123, 2005.

[11] RSA SecurID. http://www.rsa.com/node.aspx?id=1156.

[12] J. H. Saltzer and M. D. Schroeder. The protection of information in computer systems. *Communications of the ACM*, 17(7), July 1974.

[13] A. Seshadri, A. Perrig, L. van Doorn, and P. Khosla. SWAtt: Software-based attestation for embedded devices. In *Proceedings of the IEEE Symposium on Security and Privacy*, May 2004.

[14] S. Sinclair and S. W. Smith. PorKI: Making user PKI safe on machines of heterogeneous trustworthiness. In *21st Annual Computer Security Applications Conference*, pages 419–430, Los Alamitos, CA, USA, 2005.

[15] J. Yan, A. Blackwell, R. Anderson, and A. Grant. Password memorability and security: Empirical results. *IEEE Security and Privacy*, 2(5):25–31, 2004.