2017

Deceptive security based on authentication profiling

Andrew Nicholson
De Monfort University

Helge Janicke
De Montfort University

Andrew Jones
De Montfort University

Adeeb Alnajaar
De Montfort University

DOI: 10.4225/75/5a84f8fe95b4f
Nicholson, A., Janicke, H., Jones, A., & Alnajaar, A. (2017). Deceptive security based on authentication profiling. In Valli, C. (Ed.). (2017). The Proceedings of 15th Australian Information Security Management Conference, 5-6 December, 2017, Edith Cowan University, Perth, Western Australia. (pp.140-148).

DECEPTIVE SECURITY BASED ON AUTHENTICATION PROFILING

Andrew Nicholson, Helge Janicke, Andrew Jones, Adeeb Alnajaar
Cyber Technology Institute, De Montfort University, United Kingdom
heljanic@dmu.ac.uk

Abstract

Passwords are broken. Multi-factor Authentication overcomes password insecurities, but its potentials are often not realised. This article presents InSight, a system to actively identify perpetrators by deceitful adaptation of the accessible system resources using Multi-factor Authentication profiles. This approach improves authentication reliability and attributes users by computing trust scores against profiles. Based on this score, certain functionality is locked, unlocked, buffered, or redirected to a deceptive honeypot, which is used for attribution. The novelty of this approach is twofold; a profile-based multi-factor authentication approach that is combined with a gradient, deceptive honeypot.

Keywords: Authentication, Multi-Factor, Deceptive Security, Trust, Honeypot

INTRODUCTION

This paper addresses an aspect of multi-factor-authentication in combining a number of behavioural indicators to provide more trust in the identity of the user. Multi-factor authentication, especially when based on behavioural metrics such as key-stroke recognition provide a level of trust in the identity of the user. Typically such systems use a threshold to either authenticate or deny access to a user. The novelty of the approach is that instead of only authenticating a user, the system is dynamically configuring a deceptive honeypot to include additional attribution techniques to ascertain the identity of the actual user. The result is the InSight system that adjusts dynamically the deployment of deceptive features and attribution techniques based on the trust-level established through the user's interaction with the system.

Weaknesses have long been identified in traditional authentication systems based on username and password (Adams 1999) (Ives 2004) (Schaffer 2011). A typical person is capable of remembering 4 to 5 passwords, however, research by (Sasse 2011) shows that at work a person is likely to need in the region of 15 to 16 different passwords. This creates bad practices, such as writing passwords on post-it notes and reusing the same password between systems. When one system is cracked, all others topple over like a line of dominoes (Ives 2004).

Worse still, people often choose memorable personal passwords, such as their date of birth or the name of their first pet. This information is often shared on social networking websites or can be obtained through social engineering. On the other hand, when people do choose complex passwords, they often are difficult to remember and are consequently recorded in notebooks or on yellow stickers. Even complex passwords remain in many cases trivial to brute force attacks.

Security breaches at large organisations have, in the past, led to password credentials being posted publicly to websites such as pastebin.com or through bittorrent sharing websites. An analysis of these credentials often shows easy to crack login credentials. For example, in one corpus of 32 million credentials, it was found that the average password length was between 6 and 9 characters. Further analysis of the data set showed that the most popular password was ``123456" with 290731 unique occurrences[1].

Adversaries typically use three well known methods to compromise passwords; dictionary attacks, brute force attacks and rainbow tables. Simple passwords such as 'Password' and 'princess' can be compromised by an adversary using a dictionary based attack. A dictionary file contains a list of words which is used to sequentially attempt authentication. Dictionary attacks are fast; using a typical home desktop computer, dictionary words are trivially compromised in a matter of seconds. Dictionaries even exist that replace characters with common substituted character variations (e.g. `p4s\$w0rd`).

[1] http://www.imperva.com/docs/WP_Consumer_Password_Worst_Practices.pdf

Alternatively an adversary may perform a brute force attack to compromise a password by using all possible permutations of available characters. This is a slower attack, but has the potential to crack any password. If an adversary knows something about the password (e.g. it is greater than 8 characters and it ends with a number) then this can be incorporated into the brute force attempt to reduce the effective search space and time required to compromise the password. This is also known as a mask attack.

Adversaries can use rainbow tables, a pre-computed file containing all possible one-way hashes for a given dictionary. However, it is considered unrealistic to generate rainbow tables for all possible salt + password combinations. Recent advancements in graphical processing units (GPU) shows us that passwords can be cracked at a rate of 3.3 billion passwords per second. A common practice in the security of user passwords stored in databases is to use a cryptographic salt. This should be a complex, difficult to guess, string that is added to the user's password before it is encrypted. Therefore a standard dictionary attack or use of rainbow tables is likely to fail since the encrypted one-way hash is a representation of the salt + the password. However, if the adversary knows the salt then they can generate a rainbow table which includes the salt.

A solution to this problem that is in widespread use is multi-factor authentication, which combines the username and password with an additional authentication metric such as a biometric fingerprint scan. The three widely accepted authentication principles for the identification of a user are shown in Table 1.

Table 1: Widely accepted Authentication Principles

Authentication Principle	*Example*
i) Something the User has	Authentication Token, RFID Card
ii) Something the User knows	Password, PIN, Passphrase
iii) Something the User is	Fingerprint, Iris Scan, Voice Pattern

Multi-Factor Authentication

Multi-factor authentication typically uses a combination of authentication principles to establish a user's identity (Nag 2015). For example, the credit card payment system (Kumar 2008) with biometric authentication employs fingerprint verification with a credit card, combining principles i) the card and iii) the fingerprint. Such an approach requires the installation of additional equipment, increasing the cost of the approach and preventing a more widespread adoption. The use of additional devices such as fingerprint readers typically adds to the time taken for authentication which also affects the user acceptance of the technology. Depending on the technology used, fingerprints can also be spoofed (Ihmaidi 2006). Worse, usable fingerprints may readily be available on the credit card itself. Most current approaches to multi-factor authentication are expensive, difficult to deploy and directly affect the usability of the system as they prolong the authentication process (Naji 2011).

A well recognised alternative to fingerprints are keystroke dynamics which can be used as behavioural biometrics for users. It is an analysis technique in which the typing behaviour of users whilst inputting through a keyboard input is monitored (Oppliger 2011). However, if a keystroke is not combined with particular keystroke keys such as the password, it is insufficient to be an objective authentication factor (Teh 2010). Another challenge in this area is that the way users type is very much dependent on the devices they use to enter their credentials. This led to problems in accuracy with this authentication approach when users are able to use a variety of hardware configurations, such as laptops, tablet PCs or smart-phones for data entry. Indeed the use of hardware for authentication has been used since the 1980s. The idea was that users register their devices, e.g. based on their MAC address, so that the devices are authenticated rather than their users.

The authentication module of InSight combines hardware authentication with keystroke recognition to overcome some of their respective problems. To be compatible with traditional authentication approaches InSight extends a simple password mechanism with additional profiling techniques to create a form of multi-factor authentication that is based on hardware and behaviour profiles. Both of these additional factors do not require the user to memorise or otherwise keep any additional secret information.

This also means that InSight does not require special devices to be deployed to end-users, avoiding the impact of additional authentication procedures on usability. InSight integrates profiling information with the established username/password authentication thus discriminating the valid use of password credentials against misuse by establishing a level of trust in the authenticity of the user. This level of trust is then used to drive a deceptive back-end for attribution purposes that adapts the level of deception corresponding to the trust in the authenticity

of the user. The authentication in InSight is based on the trust in the multi-factor assessment, whereas other approaches (Koved 2015) take a risk-based approach.

InSight determines the user behaviour with respect to entering usernames and passwords in correlation with users' hardware. First InSight determines the assumed identity of the user through normal password authentication, gathering in additional information about the hardware configuration and the keystroke behaviour of the user when typing the username and password. The username is then matched against the hardware configurations stored in an associated hardware profile. This compares the ownership and usage patterns of the hardware. Based on the hardware profile the login procedure discriminates between keystroke profiles against which the current login request is evaluated.

Deceptive Security with Honeypots

Honeypots are specially crafted systems that lure adversaries by imitating vulnerable systems, services and software. Honeypots monitor the interaction between the adversary and the system so that the collected data can be analysed by an investigator or automated process. Honeypots are capable of misleading adversaries into revealing information about themselves, by e.g. inadvertently revealing their preferred tools and techniques, coding mistakes and hours of operation. Honeypots are often classified by their fidelity; low or high. Low interaction honeypots, such as Dionaea or Nepenthes, generally simulate a single service and are effective when facing automated scripts such as worms. They are easy to deploy and manage, however they are quickly identified as honeypots by human adversaries, since they offer only limited interactivity. High interaction honeypots are fully fledged operating systems hosted on physical equipment or in a virtual environment. They require high levels of human monitoring and there is an increased risk that they may be compromised and used by an adversary e.g. as a node in a botnet of machines. However, they offer much higher levels of fidelity, such that there is less chance of them being identified as a honeypot. Tools such as Sebek[2] are used to monitor high interaction honeypots and use rootkit techniques to hide deep inside the operating system to avoid detection.

A recent trend in honeypot research has been to combine honeypots with other security technologies. Honeypots have been stitched together with intrusion detection systems, firewalls and host-based security technologies such as anti-virus. While network-based services were once the only strand of honeypot research; a wider variety in the form of wireless, USB, bluetooth, client-based honeypots and honeypots in non-IP networks such as industrial control systems characterise this area of research.

InSight uses the level of trust that is established by the authentication module, to adapt the underlying system's functionality. This approach has the result that the user, or adversary, cannot distinguish the real system functionality from the one that is provided by InSight's adaptive honeypot. The novelty here is to not simply to redirect the user to a honeypot, but to have a continuous gradient between the real system and functions that are not influencing the system with direct and immediate effect.

InSight's Honeypot Trust Model

Deceptive technologies are an unusual but suitable partner to multi-factor authentication systems. This partnership can be justified by challengeing one of the primary principles of authentication: *Upon failing any part of an authentication challenge, a user should be denied access to the system.*

Traditionally, when failing one or more checks in a multi-factor authentication, the actor is not authenticated and consequently denied access to the resources. InSight relaxes this authentication principle using indicators of malice (IOM) to determine a trust-level in the actor's authenticity using a weighted average of the IOMs derived from the Hardware and the Behaviour profile. When an actor attempts to login using a legitimate username/password combination but the device signature and/or the behavioural biometric does not match the profile in the database, then these are considered IOM. Currently InSight only supports Hardware Profiles and Keystroke Recognition as a biometric, but this can easily be extended to include other IOMs as indicated in Figure 1.

As a result of the computed trust score, the actor is placed within a system that employs a gradient from high trust, in which all system functionality is enabled, and with few characteristics of a honeypot, to low trust, in which actions are buffered, and can be rolled-back as well as profiling, traceback and other deceptions enabled.

Figure 1 shows how the honeypot functionality increases as trust is decreased. Trust is decreased as users trigger IOMs indicating that they are potential adversaries that succeeded in obtaining username/password pairs, but

[2] https://projects.honeynet.org/sebek/

without matching the behaviour profiles adequately. To equal proportions, access to the functionality of the real system is minimised as trust is decreased. The partnering of multi-factor authentication with deceptive security in InSight, results in the authentication matrix shown in Table 2.

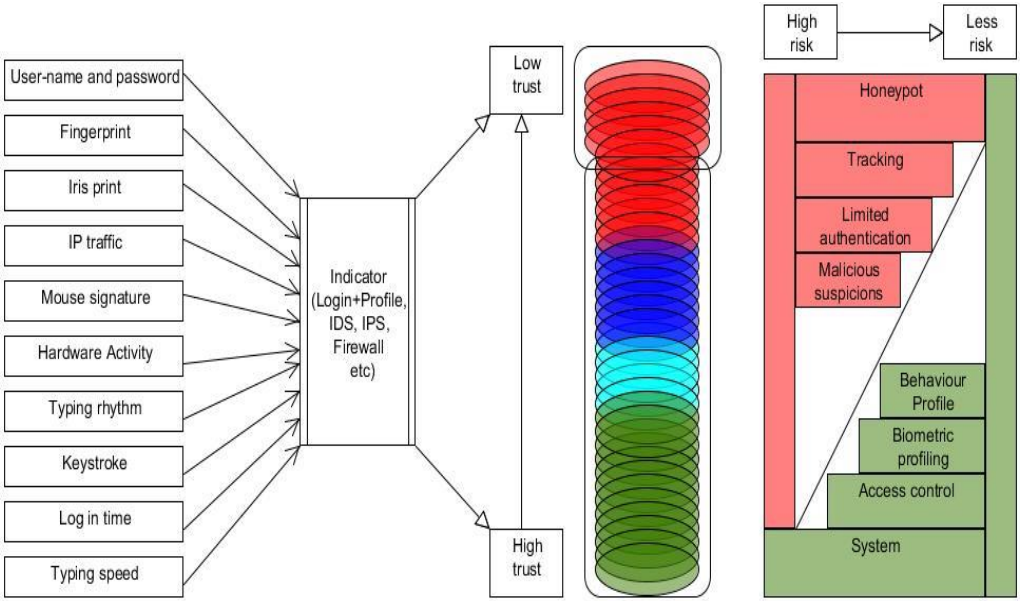

Fig 1: Honeypot Trust Model

Table 2: Multi-factor Authentication with and without InSight

Username Password	Hardware Profile	Behaviour Profile	Without InSight	With InSight
✓	✓	✓	Success, user logs in	Success, user logs in
✓	✓	✗	Denied	Partial success, user logs in to InSight system with a computed trust score
✓	✗	✓	Denied	Partial success, user logs in to InSight system with a computed trust score
✗	✓	✓	Denied	Denied
✗	✗	✗	Denied	Denied

Experiments

Our experiment is based on an online (e-)banking system. E-banking allows for customers to conduct online financial matters, such as transactions and mortgage applications through the Internet. The mid-90s saw the birth of online banking; now its use is widespread. Recent studies identified that 30% of the UK population use online banking. Banks benefit from cost savings as their branches and telephone call centres receive less interaction, customers benefit from convenient and fast access to key financial actions such as transferring funds and paying bills. However, financial fraud is a pressing issue and solutions have been proposed by security researchers (Oppliger 2011). The IC3 reported losses of hundreds of millions of dollars due to online account takeovers and unauthorised funds transfers between 2005-2009.[3]

User authentication plays a critical role as customers typically login using a web or smartphone application. Banks have adopted multi-factor authentication on a widespread scale. One-time PIN hardware devices have been sent to millions of customers free of charge by major UK banks.

In the experiments described in this paper the login server belongs to a fictional high street bank that employs numerous security methods to identify cybercriminals. The experiment is based on synthetic datasets, that were obtained through a 12 student volunteers logging into the system and processing pre-defined tasks. In subsequent experiments this will be expanded to use field data and a real-world application. The prototype experiment s

[3] http://www.ic3.gov/media/2010/100312.aspx

assume that both legitimate and illegitimate users have access to valid username and password credentials. The illegitimate users are assumed to have acquired the credentials via, e.g. compromised machines, by installing a key-logger or rootkit.

InSight grants access to the banking resources provided that they have correctly entered the username and password credentials. Unlike traditional multi-factor authentication, the user is also granted access when the hardware and behavioural profiles are not matched. This results in a lower trust score, representing an Indicator of Malice. Insight will consequently monitor their activities and deploy attribution mechanisms such as honeypots. To show the practicality of the Insight approach the honeypot features allow for delayed execution mechanisms and confirmations through independent channels. For example when money is moved from one account to another, InSight lets them appear to the illegitimate user as having been processed legitimately, thus increasing potentially the interaction with the honeypot. Two alternatives to achieve this behaviour are:

1. Carry out and then later roll-back transactions - Any actions that take place on a compromised account would be rolled back to their original state. In our e-banking scenario, this technique is problematic when funds are transferred to external accounts.

2. Buffer transactions - Actions are buffered for a given amount of time or until a certain condition is met. An example of such is that enough attribution data has been collected. Based on the user's trust-score adjusted over a period of time, Insight can then choose to either process or drop the action. This approach will produce a lag in transactions which a nefarious user may become suspicious of. This approach is deemed to be acceptable as banking transactions already have similar systems in place, causing a similar lag.

Figure 2 shows the simplified banking website from the observation point of a logged in user.

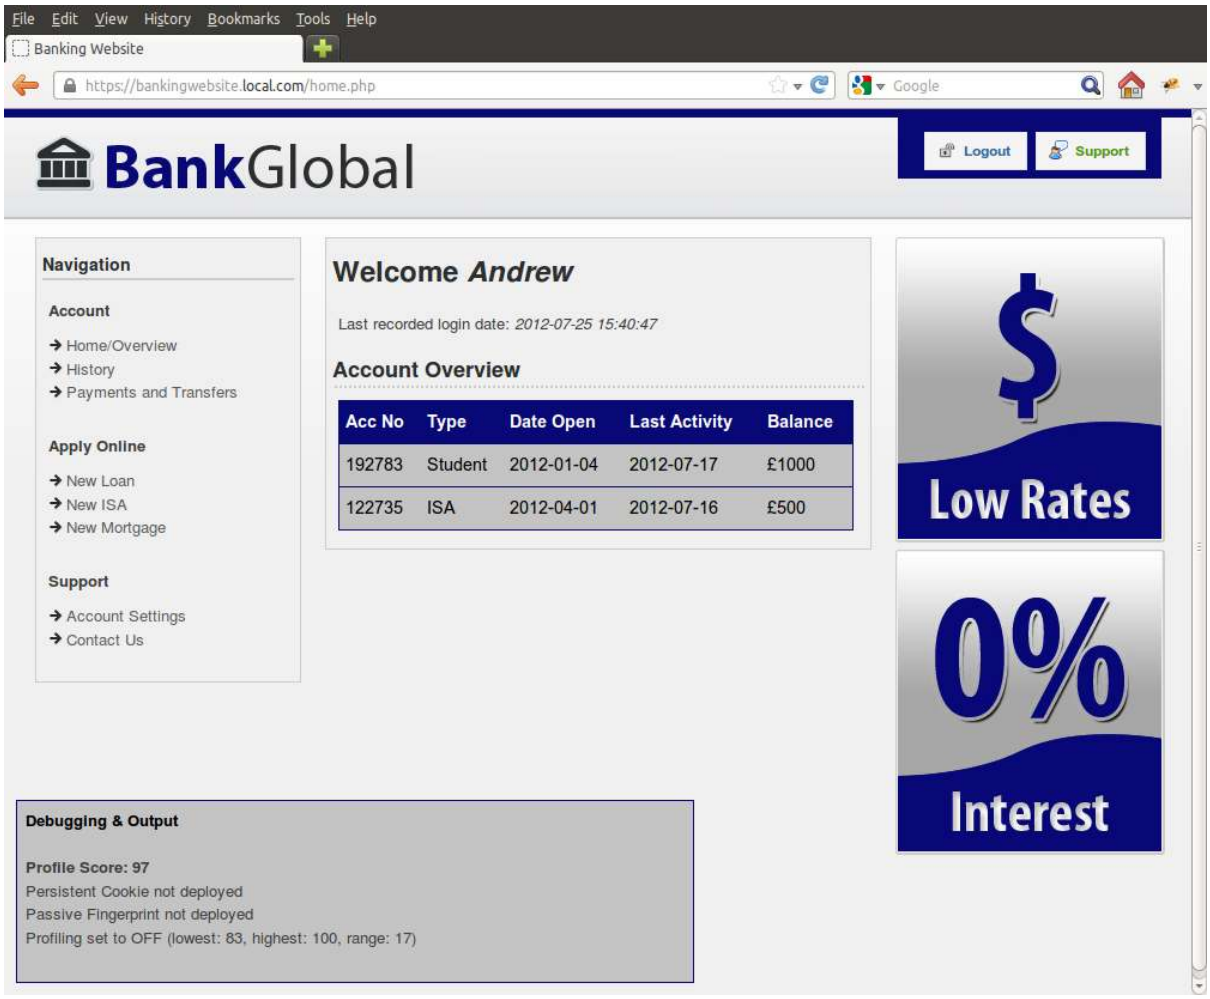

Figure 2: Simplified Banking Website for Experiment

The experiment consists of a number of realistic use cases that represent legitimate and illegitimate users who have obtained valid username/password combinations. The use cases for the experiment areas follows:

Legitimate User:

- Prerequisites: none
- Interaction Flow: Logs In Using Correct Credentials, Common machine and Correct Behavioural Biometric → InSight assesses credentials, common machine and correct behavioural biometric. User is logged in, → User is not monitored, makes required transactions → InSight accepts → User logs out.

Illegitimate User 1 (Local Nefarious User)

- Prerequisites: Acquired correct credentials by installing a keylogger onto a local shared machine (e.g. library). The illegitimate user uses the same machine to login with the stolen credentials. This user has physical access to the machine and therefore has a correct common machine and thus correct device signature.
- Interaction Flow: Logs In Using Correct Credentials, Common machine and Incorrect Behavioural Biometric → InSight assesses credentials are correct, common machine is correct and biometric is incorrect. User is logged in, → User is monitored, makes required transactions → InSight buffers → User logs out.

Illegitimate User 2 (Remote Nefarious User)

- Prerequisites: Acquired correct credentials by compromising a legitimate user's remote machine. The illegitimate user compromised the machine with a phishing email and installed a rootkit which contains a keylogger, which is able to record all interactions (e.g. keystrokes) and covertly send them to a remote server which the illegitimate user controls.
- Interaction Flow: Logs In Using Correct Credentials, Non-Common machine and Incorrect Behavioural Biometric → InSight assesses credentials → correct, common machine → incorrect and biometric → incorrect. User is logged in, → User is monitored, makes required transactions → InSight buffers → User logs out.

RESULTS

We find that within a controlled environment InSight operates as expected. The login screen provides controlled test accounts for each of the three outlined users (legitimate, local nefarious and remote nefarious). Figure 3 showed a successful login to the system by a legitimate user. The debugging and output panel shown in the bottom left corner shows the current trust-level of the user and the deployed deception mechanisms. It is only present during debugging mode and is described further below.

Figure 3: Simplified Backing Website, Experimental Results

Figure 3 shows the web-based debugging and output panel that were used during experiments. The honeypot in this prototype features three features used for attribution: i) persistent cookies; ii) passive fingerprinting; iii) behaviour interaction profiling. The profile score is computed through the use of time-of-use and place-of-use profile combined with a key-stroke metrics to create a biometric. 3A shows the results of a legitimate user login. In this case the output states that no honeypot functionality is deployed. 3B shows the results of the first illegitimate user login. In this case a persistent cookie has been deployed to track the user, however, passive fingerprinting has not been deployed. At this level and below transactions are buffered. Finally 3C shows a low profile score, using correct credentials, but the hardware profile and behavioural characteristics have failed. In this case the user has been assigned a low profile score and all methods of profiling and tracking are enabled.

Figure 3 also demonstrates that the range of the profiling technique can be defined by the operator. This e-banking scenario defined three distinct categories and the window for a trusted user is small; the trust score must be between 83 and 100. For an e-banking website this seems to be appropriate. However, adequate testing and base-lining using real customer data would help to identify a suitable range to minimise false positives and false negatives. The categories and ranges are customisable.

DISCUSSION

The previous section showed that the InSight approach performs well and delivered promising results. This section critiques the approach with respect to a real world implementation. The following are considered: i) costs, ii) false positives, iii) human factors, vi) performance factors and compromise.

Costs

Current multi-factor authentication systems require physical devices, e.g. one-time password (OTP). Newer solutions place the OTP device within the payment card itself, which is convenient for the user, but simply shift the costs into the production of the card.

Other systems such as iris scanners, fingerprint readers and RFID readers also require physical hardware to be present that is unlikely to be owned by a home user. When considering online banking, a high-street bank would need to send millions of these devices to home-users. To place this cost into perspective, personal USB device fingerprint readers currently range from GBP30 to GBP100, while eye scanners are not available for home-use. Eye scanners, such as retina scanners are costly and are invasive to individuals. The use of these devices is common in high security facilities but has not transcended to home use.

The multi-factor authentication system presented in this work does not require any new hardware to be present and therefore avoids the associated costs. Instead, the approach uses something the user already has across multiple devices; a keyboard and a device signature. So a software solution would need to be developed which implements the approach and is available on multiple device platforms.

With regard to the deceptive system, the maintenance costs can be better understood by examining the honeypots. A critical aspect is that honeypots should be maintained by a skilled team; the higher the interaction level the greater the maintenance requirements. The honeypot system for this approach need not have high interactivity, since in fact the user has full use of functionality, but their requests are buffered. The costs of honeypot maintenance may be a barrier to entry for smaller organisations. The key is the integration into the organisations' business processes. If a user is only partially trusted, the additional delay or additional effort in establishing his/her identity in the case of high-value transactions is allowing to dynamically adjust the risks to the organisation against experience of the user.

False Positives

If a user is having a stressful day and their behaviour profile is unusual, they may slip into the honeypot side of the gradient; this is a false positive. This may result in their activities being monitored and their requests, such as banking transactions, buffered. Depending on the system, this may be highly inconvenient. In fact, this result is not problematic two reasons. One, it is anticipated that the user would be contacted by an additional channel (e.g. telephone), as is already standard practice for unusual banking transactions. Two, this increased monitoring is not wasted; it can be used as part of a feedback loop which further improves the system and the novel approach, creating a better baseline of the correct user. With regard to false positives, the approach is similar to intrusion detection systems. Tuning is required and depending on the system in question, an appropriate number of false positives and false negatives should be known from the offset.

Human Factors

The multi-factor authentication technique offers significant benefits in the real world. Humans and passwords are not a particularly good combination (Kovet 2015) (Nag 2015) (Charab 2007) (Kumar 2008). Using InSight the authentication technique is based on a behavioural pattern that is difficult to replicate and is inherent to the individual. Also, while behavioural patterns do change slowly over time, such as a user becoming adept at touch typing, the algorithm can actually account for this.

Performance Factors and Compromise

In our approach there are opportunities for compromise at the client and at the server. At the client side a nefarious user with access to a compromised machine could potentially record a user's keyboard timing data and their hardware details. These could be replayed along with the correct username and password.

The InSight server collects a corpus of behavioural data. If the system is compromised then the unique behavioural records could be exposed. Therefore it is of importance that the system is hardened to deter and prevent compromise. It is also important that behavioural data is encrypted when stored.

Regarding performance of a system that were to implement such an approach, it is likely that there would be an increase in processing and that this would increase the further into the honeypot the user is. Therefore, if the adversary was aware of the system, they might be able to identify the depth at which they are located based on timing data. One solution to this problem is to normalise output to a pre-determined and acceptable speed or duration and add randomness so that output is not 'too similar'.

CONCLUSION

This paper discussed the inherent weaknesses in single-factor authentication systems and current multi-factor authentication systems. InSight was presented and demonstrated to be a novel system that uses behavioural aspects and hardware profiles for multi-factor authentication to compute a trust score. This score is used to determine the users position in a deceptive environment. Insight offers the attractive capability of being able to create a corpus of adversary activities. While the experiment demonstrated the capability within a simulated e-banking environment, the approach could be used in any other system that requires authentication. This data is useful for detecting ongoing attacks using signature matching and is especially useful for attribution, traceback and profiling purposes.

REFERENCES

Adams A. and Sasse M. A. (1999). Users are not the enemy. Communications of the ACM, 42(12): 40–46.

Chakrabarti S. and Singbal M. (2007). Password-based authentication: Preventing dictionary attacks. Computer, 40(6):68–74.

Ihmaidi H., AlJaber A, and Hudaib A. (2006). Securing online shopping using biometric personal authentication and steganography. In Information and Communication Technologies, pages 233–238,.

Ives B., Walsh K. R., and Schneider H. (2004). The domino effect of password reuse. Communications of the ACM, 47(4):75–78.

Koved L. (2015) Usable multi-factor authentication and risk-based authorization. Technical report, DTIC Document.

Kumar D and Kwon D. (2008). A survey on biometric fingerprints: The cardless payment system. In Biometrics and Security Technologies, page 16.

Nag A. K., Roy A., and Dasgupta D (2015). An adaptive approach towards the selection of multi-factor authentication. In Computational Intelligence, 2015 IEEE Symposium Series on, pages 463–472. IEEE.

Naji A. W. (2011). Security improvement of credit card online purchasing system. Scientific Research and Essays, 6(16):3357–3370.

Obaidat M. S. and Sadoun B. (2008). Keystroke dynamics based authentication. Biometrics, pages 213–229.

Oppliger R, Rytz R, and Holderegger T (2009). Internet banking: Client-side attacks and protection mechanisms. Computer, 42(6):27–33.

Sasse M. A., Brostoff S., and Weirich D. (2001). Transforming the ‘weakest link’—a human/computer interaction approach to usable and effective security. BT technology journal, 19(3):122–131.

Schaffer K (2011). Are password requirements too difficult? Computer, pages 90–92.

The P.S. (2010). Keystroke dynamics in password authentication enhancement. Expert Systems with Applications, 37(12):8618–8627.