EXHIBIT 20

Incremental Security in Open, Untrusted Networks

Andrew Hutchison, Marc Welz
Department of Computer Science
University of Cape Town
Rondebosch, 7701 Republic of South Africa
{hutch,mwelz}@cs.uct.ac.za

Abstract

In this paper we identify a number of security problems encountered in open, untrusted networks and motivate why some of these problems are going to remain with us for the foreseeable future. In order to reduce system vulnerability in such environments, we suggest that network services should provide a second line of defense to catch those attackers who are not excluded by the first line — the conventional signon process. Part of this fallback position could adapt anomaly detection (a concept borrowed from conventional network intrusion detection systems) to provide a means of gradually and continuously authenticating users and modulating their access rights accordingly.

1 Introduction

Computer network connectivity costs are decreasing for the end user. At the same time it is becoming possible to access computer networks from an ever increasing variety of platforms such as cellular telephones, internet kiosks and pagers. The combination of these two trends means that unsophisticated users will become an ever increasing fraction of the online population.

We shall refer to such cheap, ubiquitous networks as *commodity networks*.

Users of such a networks (*subjects* in this context) will have to be authenticated and granted access rights to resources (referred to as *objects*). There are a number of challenges associated with this process:

- Authentication has to be reasonably simple and non-intrusive.
- Subjects are naïve and thus can't be relied on to follow good security procedures.
- It may be difficult or impossible to verify the identity of a subject.
- There exists a well-established and experienced intruder population.

This paper will describe these problems in greater detail and describes an approach which may be used as a second line of defense in such a hostile environment.

Our approach attempts to incorporate the anomaly detection capabilities typically only found in network intrusion detection systems (see [1] for a example of a research system or [2] for an overview of commercial ones) and make them an integral part of an application, where anomaly detection may not only be used to provide a continuous and progressive authentication mechanism, but also a means to constrain the available actions to those needed and actually used.

2 Security Challenges in Open, Untrusted Networks

2.1 Simple, Inexpensive Authentication

A requirement of a consumer network infrastructure is that authentication should be reasonably simple and inexpensive. For example, it is unlikely that ISPs will require that subscribers install retina scanners (at least at current prices) in order to access the internet from home.

Another example of ease and convenience taking precedence over security is that passwords for dialup accounts are often stored in plaintext on the local machine and changed infrequently if ever.

It appears unlikely that these trends will be reversed anytime soon — the computer industry has created the expectation that computers should be simple and easy to use, while it is probably going to be difficult to persuade the commodity PC hardware industry to add expensive authentication devices to home PCs.

0-7695-0468-X/99 $10.00 © 1999 IEEE

2.2 Naïve User Population

Despite valiant efforts by educators and support personnel, computer users still do write passwords on post-it notes stuck to their monitor. It seems unlikely that this will change — more and more people will use computers as a mere tool and won't have an interest in computers themselves.

2.3 Unverifiable Identity

In a number of situations it is difficult to associate an online user with a real person or organization. For example users of services such as prepaid cellular telephony have, for all intents and purposes, no identity. Unless the user of such a telephone chooses to tell you, there is no reasonable way of establishing his or her name.

In some situations it may be possible to trace the airtime purchase to a credit card, but requiring that prepaid cellular phones are only purchased with credit cards is not practical. To illustrate this point: In South Africa prepaid cellphones were introduced to make wireless communications available to those who would not qualify for credit. Their introduction has been credited with a significant growth in the number of South African GSM telephone users and some of these new users are reported never to have opened a bank account.

2.4 Established Intruder Population

System crackers are a part of the Internet. While a large proportion of crackers are amateurs who merely use existing cracking tools, there does exist a category of cracker who undeniably is able to mount complex attacks.

While the classical cracker is portrayed as an individual who breaks into systems for the intellectual challenge, it would seem reasonable to assume that a number of crackers are in the service of intelligence agencies, both military and commercial. Such crackers are likely to be experienced and motivated enough to keep abreast of the newest security developments.

2.5 Fundamental Security Problems

The above description is intended to show that it is difficult to secure an object in a commodity network — vulnerabilities exist at any point between it and the subject.

It might be argued that today's networks were never designed to resist determined attackers and that the next generation should be more secure. Said next generation networks are supposed to employ strong cryptographic methods, smart cards and biometrics to exclude intruders and impostors.

And while we hope that future networks will be more secure, it seems unwise to believe that all vulnerabilities will go away: Cryptographic channels might contain trapdoors and will reduce the efficacy of network intrusion detection systems or virus scanners. Biometric credentials are difficult to revoke if ever compromised. Smart cards can be stolen and don't necessarily map to an identifiable subject — users of prepaid GSM phones are still difficult to trace, despite being accompanied by smart cards.

Apart from criticisms of particular technologies, there exist two more fundamental problems:

For one it is very difficult and expensive to construct a truly secure system — given the pressure to deliver a new network service to the market as fast as possible and at the lowest cost, it is probable that security issues will not receive any more attention than they receive currently.

But even if it were easy to construct a secure network, it is still unclear if such a system is desirable: A network where each subject can be identified and mapped to a known real-world entity would offer no privacy to its users. There already exist concerns that current networks record too much information about their users: For example, rash USENET posts have come to haunt their authors at job interviews. If these trends continue reporters are likely to quiz a future presidential candidate about the web sites he visited as teenager.

Put simply, a number of real world activities (such as cash payments) are anonymous and without permanent record. If these activities are to have electronic equivalents, then some form of anonymity has to be possible. In other words there is a tradeoff between the accountability and the privacy of subjects in a network. If it is desirable to grant subjects some degree of privacy then there exits the opportunity for hostile subjects to launch attacks.

3 A Second Line of Defense

The above suggests that hostile subjects are always likely to probe objects on a commodity network, and that the owners of such an object may not be able to do very much about this — the attacker may use an anonymous service, use a stolen identity, launch an attack from a compromised intermediate or be based somewhere where the victim has no legal recourse.

Since it does not appear feasible to exclude hostile or naïve subjects from a commodity network, we propose that a second line of defense be made a standard component of distributed applications.

Where the first line of defense includes conventional subject authentication (via password, smartcard or fingerprint), the second line uses an alternative means to identify a user.

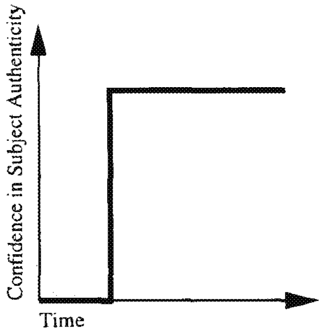

Figure 1. Conventional authentication (binary value)

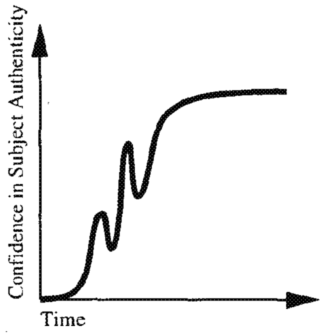

Figure 2. Progressive authentication (fuzzy value)

A first line of defense exists in most distributed systems: subjects usually have to pass an initial authentication phase. Once a subject has passed (or bypassed) this phase the subject is granted access to a set of objects.

The point to note is that the above security measure consists of an initial phase where after no security checks are performed.

We suggest that the second line use the actions of a subject as a way of verifying the identity of a user. This has the advantage that the authentication module is in operation for as long as the subject is accessing the object, also this security measure can be implemented entirely on the side of the object, and requires no co-operation of or trust in the the subject. Furthermore, such a system would no longer restrict confidence in the user authenticity to a binary value (yes, no), instead it would be possible to have a progressive gradation, and be able to adjust access rights accordingly: For example, host Bilbo has a confidence factor of 0.95 that Mr Jones' account is being used by its rightful owner since

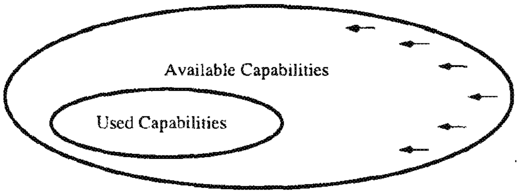

Figure 3. Gradual reduction of capabilities to those exercised

Mr Jones usually logs in at 7:20 and first checks his mail before checking his diary. A factor of below 0.6 would restrict the user of Mr Jones' account to checking his mail, and a factor of below 0.2 might page the system administrator.

An object which would implement such a second line of defense would be equipped with the following two components:

- A module which profiles subject activity in order to establish usage patterns and trends. Where anomalous behaviour patterns emerge, the system may flag alerts or disable a service. Where volumes of data are too large or where privacy issues prevent full logging, it seems worthwhile to investigate inscrutable pattern matching techniques such as neural networks or genetic algorithms since these can be thought of as maintaining only a digest of past user behaviour, and can thus not be used to reconstruct an exact record of past user behaviour.

- A component which establishes what services are not being used by a particular subject (possibly using the module explained in the previous paragraph), with an option to temporarily or permanently disable such services. For example, a given user might only use a home banking service to examine her current balance. The proposed component might then notify the user that her ability to initiate transfers would be disabled unless this component received verified instructions to the contrary. Such a component would protect unsophisticated users who do not make full use of a given service. The component can be thought of as a way of automating the principle of least necessary privilege, since the component would gradually restrict the users rights to only those privileges needed and exercised.

We do note that these ideas are not new (see [3]) for an example of a host-based IDS, while [4, 5] use an immune systems metaphor) — anomaly detection has been part of network intrusion detections systems for some time. However, the use of anomaly detection modules as an integral part of an application does not yet seem to have been explored fully.

As mentioned previously, we are particularly interested in investigating how the complement of anomaly detection (ie detecting normal behaviour) can be used to provide a continuous and progressive means of authenticating a user (one might call this fuzzy logic for authentication, since confidence in user authenticity ceases to be a binary value), and how this confidence value can be used to modulate the access rights of the subject. Our second, related, area of interest involves the use of an anomaly detection/profiling system to determine the set of actions typically performed by a subject (versus the set of possible actions), and reducing the set of possible actions to those used (one might refer to this as the *If you don't use it, you loose it* principle). This would offer an automatic way of implementing a least privilege policy.

We anticipate that anomaly detection will coupled ever more closer to applications or services — apart from the above-mentioned possibilities, a tighter coupling would also offer a number of other advantages, including a reduced development effort (it would require less effort to keep the two synchronized) and easier access to application state (this will become increasingly important if network traffic is encrypted, since encrypted traffic would degrade the efficacy of a conventional network intrusion detection system significantly).

4 Applications and Limitations

Our proposed second line of defense is likely to be most effective in situations where authorized subjects perform a small set of tasks — abnormalities are recognized more easily under these circumstances. As it turns out, naïve users, the largest fraction of commodity networks users, do fall into this category — these users typically only use a limited subset of a particular application. By automatically disabling, or at least monitoring the use of more sophisticated features, it should be able to detect a number of abuses. For example, a naïve user is unlikely to take advantage of the macro capabilities of a word processor, thus the sudden use of sophisticated macros might be indicative of a macro virus infection and should thus trigger an alert.

The corollary of this observation is that an anomaly detection system is of lesser use where subjects are sophisticated and perform a large set of complex operations. While this does present a problem, it is worth noting that sophisticated (as opposed to naïve) users are more likely to follow sensible security procedures (eg: selected complex passwords, memorize passwords instead of writing them down, et cetera) and are thus, ceteris paribus, less likely to fall victim to an attack.

5 Conclusion

System crackers are likely remain a threat to commodity networks. Protection of such networks is complicated by the fact that their users are unreliable — most lack the knowledge or motivation to follow a reasonable security policy. For this reason it seems prudent to augment a conventional authentication component (based on an initial signon with password, biometric or key) with a user profiling or anomaly detection module which allows the system to verify the authenticity of a user throughout a session and adjust the users access rights, both on a per session basis (as a function of how confident the system is of the user's authenticity) and in a the long term (where access is gradually restricted to those functions actually used).

References

[1] J. M. J. Bonifacio, E. S. Moreira, A. M. Cansian, and A. C. P. L. F. Carvalho. An adaptive intrusion detection system using neural networks. In *Global IT Security*, pages 416–428, September 1998.

[2] T. Escamilla. *Intrusion Detection*. John Wiley and Sons, 1998.

[3] T. Lane and C. E. Brodley. Temporal sequence learning and data reduction for anomaly detection. In *ACM Conference on Computer and Communications Security*, pages 150–158, November 1998.

[4] C. P. Louwrens and v. S. H. Solms. Can computerized immunity be achieved, based on a biological model ? In *Global IT Security*, pages 240–250, September 1998.

[5] A. Somayaji, S. Hofmeyr, and F. S. Principles of a computer immune system. In *New Security Paradigms Workshop*, pages 75–82, September 1997.