PENETRATION TESTING PROFESSIONAL ETHICS: A CONCEPTUAL MODEL AND TAXONOMY

Justin D. Pierce[1], Ashley G. Jones[2] and Matthew J. Warren[1]

1: School of Information Systems
Deakin University
Geelong, Victoria, Australia

Email: jdpierce@deakin.edu.au

2: Information Risk Analyst
ABN AMRO
Level 27, 367 Collins Street,
Melbourne VIC AUSTRALIA

ABSTRACT

In an environment where commercial software is continually patched to correct security flaws, penetration testing can provide organisations with a realistic assessment of their security posture. Penetration testing uses the same principles as criminal hackers to penetrate corporate networks and thereby verify the presence of software vulnerabilities. Network administrators can use the results of a penetration test to correct flaws and improve overall security. The use of hacking techniques, however, raises several ethical questions that centre on the integrity of the tester to maintain professional distance and uphold the profession. This paper discusses the ethics of penetration testing and presents our conceptual model and revised taxonomy.
Keywords: Penetration testing, computer security and computer ethics.

INTRODUCTION

The adage 'knowledge is power' can be used to illustrate the gap between information security professionals and ordinary end-users. To laymen, information and especially computer security is a clandestine unknown against which they feel powerless. Bereft of security knowledge, end-users succumb to fear, uncertainty and doubt (FUD). The unbridled pace of today's business has underpinned an explosive growth in the adoption of information and communications technologies (ICTs). In an increasingly uncertain and competitive business environment, product life-cycles are shortened and their outputs pushed to market quicker in an effort to maximise profits margins: there are few industries immune to this cycle. Software companies, for example, are pressured to release applications faster and this often results in haphazard software testing practices. In fact, the underlying philosophy of the rapid application development (RAD) 'phased development' methodology centres on finishing core functionality quickly and then implementing other functionality (including security) and rigorous testing in subsequent software versions (Dennis et al., 2002). The literature shows little convergence enumerating the volume of vulnerabilities routinely discovered every week but commercially the trend remains to be playing a game of 'catch-up' in patching vulnerable software.

In a predominantly networked society these defective ICTs are transmitting vast quantities of sensitive data across relatively unsecured communications media. Moreover maturing technologies such as encryption can be used to make sensitive data useless to interceptors but it is essentially a patch in itself: the Internet was designed to survive a nuclear assault during the Cold War and not for the widespread commercial use it has adapted to today nor with security in mind. The question facing many network administrators then is 'how secure is my network?' Assessing the security of a network can be achieved using controlled hacking techniques. Penetration Testing, as it is termed, can provide security assurances to network administrators. The practice tends toward the patching trend noted but by its very nature—hacking—raises several ethical concerns. This paper seeks to discuss such ethical concerns and present a conceptual model and taxonomy to show their relationships. We start by discussing the limited body of literature before presenting a categorised set of penetration testing ethics. The paper is then concluded with directions for future research.

CONTEMPORARY PENETRATION TESTING IN THE LITERATURE

'Quality assurance and testing organizations are tasked with the broad objective of assuring that a software application fulfils its functional business requirements.' (Arkin et al., 2005, p.84) The major theme in the penetration testing literature tends toward describing how the tester should conduct their tests according to a plethora of differing methodologies and philosophies using a growing collection of labyrinthine automated tools. The commonly accepted definition of penetration testing is the '[sanctioned] illegitimate acquisition of legitimate authority.' (Geer and Harthorne, 2002, p.1; Logan and Clarkson, 2005; Thompson, 2005)

Geer and Harthorne (2002) point out that penetration testing should be considered an art rather than a science. The distinction is based on the commonly accepted limitation that penetration testing cannot prove the absence of network vulnerabilities, only the presence of them: therefore a penetration test that fails to uncover any vulnerability is not necessarily a good penetration test result (Arkin et al., 2005). Whereas science relies on the disproving of null hypotheses, penetration testing can at most be a science of insecurity as opposed to a science of security (Geer and Harthorne, 2002).

In this same vein Geer and Harthorne (2002) suggest that if a penetration test fails to uncover network vulnerabilities then it is more likely to create value for the client. The possibility for clients to misunderstand the so-called 'science of insecurity' is thus illustrated and an important question of ethics is uncovered by the bye: the chance of misrepresenting penetration testing and its potential to guarantee security. Geer and Harthorne (2002) go on (p.3) with a cynical view of penetration testing, predicting that it will evolve into more of a quality assurance regime using checklists rather than the art of discovering known and unknown vulnerabilities and providing realistic assessment of software security posture.

To the contrary, there seems to be a movement in the literature toward separating security testing from software quality assurance. In the context of bug reports and quality assurance, however, Arkin et al. (2005, p.85) suggest that

> [People] often use penetration testing as an excuse to declare victory. When a penetration test concentrates on finding and removing a small handful of bugs (and does so successfully), everyone looks good: the testers look smart for finding the problem, the builders look benevolent for acquiescing to the test, and the executives can check off the security box and get on with making money. Unfortunately, penetration testing done without any basis in security risk analysis leads to this

> situation with alarming frequency. By analogy, imagine declaring victory by finding and removing only the first one or two bugs encountered during system testing!

Haphazard penetration testing is giving way to engineering approaches emerging from the literature and will wane when best-practice standards mature and the distinction between security professionals and hackers becomes even more prominent. For example Thompson (2005) describes a novel approach to operational penetration testing where a threat model resembling a tree-like flowchart is developed. Network vulnerabilities are then tested according to reusable and evolving threat models. Further, Beznosov and Kruchten (2005, pp.49-50) describe modern tactics for assuring software quality:

> A fundamental practice in the assurance business is to keep developers and security evaluators "at arm's length" from each other so that they do not affect each other's ideas. Since security assurance must be completely neutral and objective, its practitioners and the developers should not become too closely involved except during their information gathering sessions. This leads to developers often focusing on the functional development with a "tunnel vision" that becomes quite blind to security flaws.

Recently, universities have ventured toward offering security testing courses. While still in its infancy this branch of information management is evolving into a specialised profession that will soon require undergraduate qualification like mainstream computer science. In this instance, it would seem that, academia has caught wind of industrial certification bodies such as The International Information Systems Security Certification Consortium (ISC^2), The Information Systems Audit and Control Association (ISACA), The Institute for Security and Open Methodologies (ISECOM), Cisco and Microsoft who offer specialist security certification. (The business models are slightly different for academia and the cited industrial certification bodies that exact revenue from the continued subscriptions required to maintain certification.) Nonetheless opportunity exists for academia and industry to collaborate on educating tomorrow's security professionals, yet there are several problems to iron out.

In early pedagogical case studies Tikekar (2003) and Logan and Clarkson (2005) point out that until recently, students were not required to study computer ethics to graduate. Security professionals who are not educated in computer ethics poses perhaps the fundamental question as Logan and Clarkson (2005, p.159) ask 'What happens when universities, such as Marshall University, West Virginia house the state's digital evidence lab for the state police, as well as student computer forensic training labs?' Without computer ethics curricula, graduates might be less informed as to what constitutes abuse and misuse of their skills. This scenario could extend to the possibility of al-Qaeda (and similar) recruits enrolling in such courses with the predisposed intention to later launch co-ordinated terrorist attacks on critical information infrastructure. Indeed it has been revealed that al-Qaeda used strong encryption algorithms (PGP) to hide their communications (Oz, 2006).

Penetration testing is an evolving practice and the small but growing body of literature shows that there are many arising issues that need to be addressed before it can mature fully. The literature is rich with methodologies and frameworks (see Pierce et al., 2005a; Pierce et al., 2004), but lacks longitudinal studies that prove their merits. For example, the literature lacks case studies that demonstrate how penetration testing fits into business and military continuity planning. Furthermore, although equivalent international standards such as ISO 17799-2000 are in place a formal Australian Standard for penetration testing as yet does not exist. The following section demonstrates the ethical concerns that arise from penetration testing and shows how they converge on six major themes.

THE ETHICS OF PENETRATION TESTING

The Open Source Security Testing Methodology Manual (OSSTMM) (Herzog, 2003) outlines its 'rules of engagement' which are essentially a set of rules designed to restrict unethical penetration testing practices. While the OSSTMM does not discuss the ethics of penetration testing explicitly, its scope is limited to outlining a methodology for penetration testing. We discuss the ethics of penetration testing in this section using the rules of engagement and those identified in the literature review as reference points.

The six major themes of penetration testing ethics as identified in the literature centre on integrity, which branches out to serving and protecting the client and preserving the security profession. These objectives are met by avoiding conflicts of interest, the provision of false positives and false negatives, and finally legally binding testers to their ethical obligations in the contract.

By analogy, let us imagine a profession-wide vision and mission statement: we will strive to provide security assurance consultation with professional integrity. Fulfilling the mission statement would be achieved by adopting two important and virtuous strategies: upholding the profession, and serving and protecting the client. Policies for implementing the strategies include: avoiding conflicts of interest, avoiding the provision of false positives and false negatives, and legally binding the tester's and the client's ethics. We discuss the ethical considerations in each category below.

Serve and Protect the Client and Uphold the Security Profession

Testing should not be performed without the expressed written permission of the client. Whereas in the hacking community attacks occur non-consensually, contractual arrangements must be in place to provide a degree of separation between hackers and security professionals.

There is general consensus in the literature that testers should not rely solely on automated tools but also on their skills. Notwithstanding the tester should be well-versed in computer security and know how their tool arsenal works. Tools should be tested themselves in an isolated laboratory prior to being used in production.

The use of past client's data, with or without permission, should not be used to promote the services of the tester. While it may provide a false positive to the potential client, it could also damage the reputation of the implicated organisation.

The tester should notify the client at the first instance of discovering highly vulnerable flaws as in the case of those that endanger human life. The notification should contain appropriate countermeasures to correct the flaw and minimise dangers to human life and the organisation in general.

The principle objective of penetration testing is to test security measures in a network: there is little point to testing systems known to be highly vulnerable. Testing should not commence until appropriate security has been applied to the system.

The results of social engineering tests should be delivered only in summarised and statistical format so as not to implicate individuals. The emphasis is on protecting the client and insulating unknowing

employees that might be subject to subsequent embarrassment or termination of employment as a result of the test.

The delivery of the report should be preceded by a notice of delivery. The client, upon receiving the report, should acknowledge that they are in receipt of the report. The curtesy underscores the importance of client confidentiality. The report will contain an appropriate level of detail of the tests performed, the results and the steps the client should take to improve overall network security.

False Positives and False Negatives and Conflicts of Interest

Penetration tests that fail to uncover vulnerabilities should not be passed up as free services (Geer and Harthorne, 2002; Herzog, 2003). This represents false positives by misrepresenting the penetration testing practice as an assurance of the absence of vulnerabilities.

The tester is ethically bound to serve the customer. This holds true even if it is in the best interest of the customer to engage a different testing company. In that event, the tester should not recommend any particular company so as to avoid the possibility of a conflict of interest or the perception of one.

The promoting of public hacking or trespass contests for security assurance is unethical because it implies a false security guarantee. Contests also draw unnecessary attention to the client network as a perception of a 'fair game' target will endure in the hacking community far longer than the expiration of the testing contract. When new vulnerabilities ensue the client's network may be the target of continuing non-solicited attacks.

As alluded to in the previous sub-section, clients should behave in a manner that does not encroach on the tester or interfere with a test in a manner that may alter its outcome. This includes deploying additional security during a test. As testing provides a snapshot in time of the security posture of a network, changing the security environment could lead to false positives or false negatives. Further, the client should notify only key internal personnel of the penetration test. The extent to which people are kept inside the circle is at the discretion of the client but it must be stressed that widespread knowledge of the test will alter behaviour and affect the outcome of the test thus promoting false positives or false negatives.

If white-box or in-house testing is requested, the tester should first perform black-box testing offsite. This concern draws attention to the use of internal security auditors that could develop tunnel vision (as insiders know the target network very well) and lead to false positives and false negatives. Statistically (AusCERT, 2005; CSI/FBI, 2005), however, system compromise originates in greater frequencies from within the organisation than outside the organisation. In this light the use of internal security auditors presents some merit.

Legally Binding Ethics and Other Considerations

The tester should work non-disclosure and limited liability clauses into the contract. The use of these clauses in the contract legally binds the tester to his ethical obligations. Non-disclosure is common practice, but the tester should also assume limited liability generally not exceeding the cost of the test for inadvertent damages incurred by the client as a result of negligent testing or malpractice. Therefore, it is in the best interests of the tester to have sharpened their skill base to provide the said degree of separation between hackers, script kiddies and security professionals.

Ethically, the contract should scope the tests and indicate emergency contacts as well as the IP addresses from which the tests are originating. Additionally, the contract should specify failsafe procedures such as recovering from DoS attacks.

It has been suggested that using FUD to sell penetration testing services is unethical. The OSSTMM explicates that crime facts and figures should not be used to promote security testing. Academics, however, routinely use crime statistics – such as those found in the recent AusCERT (2005) and CSI/FBI (2005) computer security surveys – to justify the problems with computer security, necessitate new research, and sell security courses at the university and practitioner levels. The question of using statistics as an ethical concern should be met with a balanced view of the need to educate end-users in security versus the instilment of FUD.

The ethical concerns presented in the paper centre on five interrelating themes: serving and protecting the client; upholding the security profession; conflicts of interest; false positives and false negatives; and, legally binding the tester to their professional ethics. A common factor, integrity, can be seen to tie these categories together as in the following figure. The tester's integrity should compel them to serve and protect the client while behaving ethically to preserve the integrity of the profession. Figure 1 illustrates our conceptual model of penetration testing ethics and how they correlate with a central theme of the integrity of penetration testers.

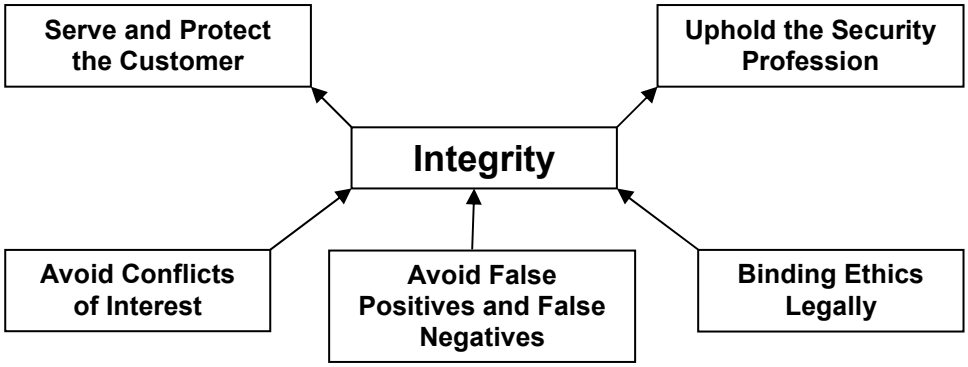

Figure 8: A Conceptual Model of Penetration Testing Ethics

Critics of an earlier version of this paper (Pierce et al. 2005b) rightly corrected that figure 1 was not, definitively, a taxonomy: it shows no specific categorisation of penetration testing ethics. Of interest in figure 1, however, is how we modelled all the ethical constructs to be related to integrity: we maintain this assertion here too. Although figure 1 shows that avoiding conflicts of interest, avoiding false positives and false negatives and legally binding ethics are prerequisites for integrity, and that integrity is a prerequisite for upholding the profession and for serving and protecting the client, we having modelled the ethics in a slightly different way in the following figure. Our revised taxonomy is a bottom-up hierarchical representation of the relationships between ethical categories as discussed in the paper. It shows upholding the profession and serving and protecting the client built upon the foundation of integrity. Additionally, avoiding conflicts of interest, avoiding false positives and false negatives and legally binding ethics are built upon a foundation of serving and protecting the client.

<table>
<tr><td rowspan="2">Uphold the Profession</td><td>Avoid Conflicts of Interest</td><td>Avoid False Positives and False Negatives</td><td>Legally Bind Ethics</td></tr>
<tr><td colspan="3">Serve and Protect the Client</td></tr>
<tr><td colspan="4">Integrity</td></tr>
</table>

Figure 2: A Taxonomy of Penetration Testing Ethics

The penetration tester should act with integrity at all times. In his endeavours, as noted, the penetration tester should strive to maintain a degree of separation between the criminal hacker and the security professional; thereby to uphold the profession. The penetration tester can be seen to be acting with integrity if they can be seen to be upholding the profession. To further clarify figure 2 suppose the penetration tester builds non-disclosure and limited liability into the testing contract, they can ergo be seen binding their ethics legally and thereby ethically serving and protecting the client. To ethically serve and protect the client is to act with integrity. Integrity is therefore the synergising foundation from which the professional ethics of penetration testing extend. If the penetration tester refuses to engage with the criminal hacking fraternity they can be seen to be using their skills for commissioned tests only and therefore upholding the profession.

CONCLUSION

We have seen that the business world is experiencing unparalleled growth due to its unbridled speed. No organisation is immune to the rapid changes in technological development that can be seen to fuel a cycle of business needs driving technological development enabling business needs. The pressures facing the business environment are also felt by software companies that resultantly cull application security testing to release the product quicker and then patch later. The problem lies in that vulnerable software is introduced to a hostile commercial environment. Penetration testing provides organisations with a means of assessing their security stance at a given moment in time. Testers and clients must behave ethically: clients so as not to alter test outcomes and testers so as to separate them from the hacking community. The ethics of penetration testing centre on integrity; serve and protect the client and uphold the security profession by behaving ethically.

The small body of literature tends toward presenting methodologies and frameworks for conducting penetration tests, but seldom integrates penetration testing into an overall business model. There seems to be confusion as to how organisations can best gauge value from the services of a penetration test. Risk analysis is a sister topic that is gaining incredible momentum in the literature and should also be integrated with penetration testing to produce an overall model of organisational security testing.

Looking to the future, the authors will investigate and propose methods of integrating penetration testing with the better established risk analysis discipline. We will also look at ways of providing pseudo-dynamic security assurance using automated approaches.

REFERENCES

Arkin, B., Stender, S. and McGraw, G. (2005) 'Software Penetration Testing', IEEE Security and Privacy, January/February 2005, pp.84-7.

Australian Computer Emergency Response Team (AusCERT) (2005) '2005 Australian Computer Crime and Security Survey' vol. 1, AusCERT, Brisbane, Australia.

Beznosov, K. and Kruchten, P. (2005) 'Towards Agile Security Assurance', Proceedings of the 2004 Workshop on New Security Paradigms, Nova Scotia, Canada.

Computer Security Institute / Federal Bureau of Investigation (CSI/FBI) (2005) 'Tenth Annual 2005 Computer Crime and Security Survey', vol. 1, CSI, San Francisco, USA.

Dennis, A., Wixom, B.H., and Tegarden, D. (2002) 'Systems Analysis and Design: An Object-Oriented Approach with UML', Wiley, New York, NY, USA.

Geer, D. and Harthorne, J. (2002) 'Penetration Testing: A Duet', Proceedings of the 18th Annual Computer Security Applications Conference (ACSAC'02).

Herzog, P. (2003) 'Open Source Security Testing Methodology Manual', ISECOM, USA.

International Standards Organization (2005) 'ISO/IEC 17799:2005 – Information Technology – Security Techniques – Code of Practice for Information Security Management', ISO, Switzerland.

Logan, P.Y. and Clarkson, A. (2005) 'Teaching Students to Hack: Curriculum Issues in Information Security', Proceedings of the 36th SIGCSE Technical Symposium on Computer Science Education, St. Louis, Missouri, USA.

Oz, E. (2006) 'Management Information Systems', 5th Ed, Thomson, Boston, MA, USA.

Pierce, J.D., Warren, M.J. and Corray, X.J. (2005A) 'In Pursuit of a Standard Penetration Testing Methodology', Journal of Information Warfare, 4(2), pp.26-39.

Pierce, J.D., Jones, A.G. and Warren, M.J. (2005B) 'A Taxonomy of Penetration Testing Ethics' in Proceedings of the 4th International Australian Institute of Computer Ethics (AiCE) Conference, September 26th 2005, Waurn Ponds, Victoria.

Pierce, J.D., Warren, M.J. and Corray, X.J. (2004) 'A Critical Review of Penetration Testing Methodologies', in Proceedings of the 5th Australian Information Warfare and Security Conference, 25-26th November 2004, Fremantle, West Australia.

Thompson, H.H. (2005) 'Application Penetration Testing' IEEE Security and Privacy, January/February 2005, pp.66-9.