



Network Security Associates

What any Organisation needs to know about a Penetration Test

Stephen Lewer
Principal Associate
Network Security Associates Ltd.

© 1999

Any organisation contemplating a penetration test should understand the serious issues surrounding the decision. The people championing the penetration test should also analyse their reasons for doing so; the introspection may be enlightening.

Does the Penetration-Testing organisation have liability insurance? Many penetration-testing teams are under qualified for the work. Some teams use hacking tools they have not adequately tested in various environments, or do not understand the implications of the techniques they have learned. This can lead to inadvertent damage or loss of information on the target systems. All penetration testing teams should carry enough liability insurance to cover the value of the information in the systems they will be attacking. If your penetration testing team uses a denial of service attack against your DNS server, your organisation could be out of action for several hours. Be certain you can recover any losses you may incur from the testing.

Does the testing team assign or request a "cut-out" in the target organisation? A cut-out is a person within the target's incident response or notification hierarchy who is fully aware of the tests scope, technique, and schedule. This person is essentially a monitor; he can intercept any excessive response by the target's incident response process and save either unnecessary work or embarrassment. For example, in one test to which the author was privy, a penetration team, operating out of a hotel in San Jose, California, was in the third day of their testing when the FBI walked through the door, guns drawn. It took several phone calls, a get-out-of-jail-free letter from the target organisation and a trip to the San Francisco FBI office to clear up the matter. The target organisation had responded to the "intrusion" by calling the FBI, who then tracked the attack back to the hotel. This is an example of not setting up an appropriate cut-out. In this example, a cut-out would have halted the response before the organisation notified law enforcement agencies. If the target organisation does not have any highly skilled technical personnel on staff, they have two additional issues to address. First, it is a reasonable conclusion that they have an inadequate security staff for the tasks they are attempting. Second, does the penetration testing team provide someone to remain on-site to monitor the test?

What instruction and training does the penetration testing team give to the cut-out? The cut-out, or monitor, needs to understand the full scope and technique

the penetration team intends to use. If the cut-out does not fully understand what is going to happen, and when, he may not respond correctly to events as they unfold during the test. The point of a penetration test is to identify vulnerabilities in the system and the surrounding processes. The team should avoid unduly stressing the target organisation's personnel or systems. The cut-out should act decisively to avoid unnecessary work or stress. For example, in one penetration test the security staff at the target organisation detected and reported probes on the target organisation's internal network. The person responsible for incident response at the target organisation reported his intention to examine, by hand, every NT domain in the organisation for further signs of intrusion or compromise. The cut-out failed, in this example, to call off the response team; as a result, one person in the response process spent 37 consecutive hours looking for signs of a nonexistent intrusion.

What information does the penetration team require at the start of the test?

The answer should be that the penetration team needs only the IP addresses of the target systems. Some penetration teams will ask for operating system versions, network diagrams, lists of response personnel, telephone directories, and, believe it or not, firewall filter rules! Hackers, naturally, will not have access to such detailed information at the start of their attacks, so to improve the validity of the penetration test; the test penetration team should not have that information either.

What systems are "fair game" for the test? Some teams want all systems to be part of the test. This is the way real hackers operate. If it is technically advantageous for a hacker to attack an insurance company's mainframe first, that is what they will do (even if this is not a very realistic example). From a management perspective, however, this may be too risky. This is, in essence, a risk decision; if the target organisation, however, is willing to risk compromising or damaging their data warehouse, they will improve the validity of their test. Pragmatically, though, some systems may need to be off limits. Be cautious if your potential penetration team does not bring up this point during your conversations.

What are the practical limits on what the penetration test will tell us? A penetration test, generally, attempts to mimic real hackers. It gives the penetration team some idea of how the organisation will detect and respond to a real attack. There are, however, severe restrictions on the validity of such testing.

Time compression. Penetration tests usually take place over the course of several weeks. Real hackers tend to be more patient with the target systems, sometimes taking months to conduct initial probes. This means that log records that would normally be lost amongst hundreds or thousands of other records are grouped together in large, obvious blocks that any fool can see at a casual glance. This makes the intrusion detection and response capabilities look relatively better than they are, because it is easier to detect a penetration test than a real attack. The amount of information that can be gathered in a given time period is the major consideration in evaluating the validity of a given penetration test. Turning a team loose on a firewall and saying, "tell me what you find in three days" is not the way to obtain realistic results. Likewise, organisations that expect real results from a few days of consulting are not going to receive a valid penetration test.

Limited set of target systems. Realistically, the penetration test will be limited to a subset of the organisation's actual network. Most large organisations have network connections to partner organisations or service vendors. In many cases, attackers can penetrate the target organisation by first penetrating one of the tangent networks. For example, one large (Blue chip!) company had a sales office in Hong Kong. The manager of the sales office was not happy with the connection to the Internet via the overseas link to the headquarters in the City of London, so he put in a direct link to the Internet. Unfortunately for the headquarters network, the manager in Hong Kong chose a flawed firewall configuration, which allowed attackers to penetrate the sales office; there was no intra-network firewall in the configuration, so once the attacker got through the sales office, he had full run of the corporate network.

Production vs. development systems. Organisations will frequently want a penetration test of a development system, rather than a production system. The thought here is that the organisation can reduce risk by "validating" the system before rolling it into production. This is actually a sound idea. Unfortunately, the organisation is very likely to make changes to the system configuration before putting it into production. These changes frequently expose the overall system to new attacks. For example, a firewall filter rule may be changed to allow ftp to a web server; if the web server's ftp implementation isn't tightly monitored and evaluated, this minor change may create a significant exposure. From a risk perspective, then, it is better to test development systems. From a validity perspective, however, it is

better to test a production system. The compromise is to either re-test after the system is in production or to enforce very strict change control procedures once the penetration test is complete.

Limits on technique. Generally, the target organisation will want to limit the attacks to non-destructive methods. For example, most teams will not use, and most organisations will not permit, DNS cache corruption attacks to exploit transitive trust relationships. Real attackers will not limit themselves to non-destructive techniques.

Missing the point. A penetration test may be restricted to attacks from the Internet in order to save time (time= money). Or perhaps the target organisation is only concerned about attacks from the Internet. Regardless, the test may miss the point. The point is to find vulnerabilities that exist in the target organisation, not necessarily to find vulnerabilities that the organisation thinks may exist.

When asked why the U.S. does not see more espionage from network intrusions, an FBI foreign counterintelligence agent said, simply, "Because there's always an easier way." It would be good for organisations to remember that. Particularly when, ironically, the FBI web site was hacked because of an unsecured internal modem. If the penetration test is overly restricted, it will likely miss significant vulnerabilities.

Logical completeness. There can be only two results from a penetration test: the penetration test team succeeded at breaking in, or they failed. If they succeeded, the conclusion is that there is at least n ways to break into the system, with n being the number of successful techniques used. If the team fails to penetrate, however, the only logical conclusion is that the system is not vulnerable to the techniques used. This says almost nothing about the overall security of the system; I contend that if t is the set of all penetration attempts in the penetration test team's inventory, and v is the set of all possible vulnerabilities, in all possible system states, then t is much smaller than v . Failure, therefore, does not mean much. In other words, information systems are finite state machines. The number of possible states is very large. Exhaustive search of the state-space is, as a result, infeasible.

What does the team do, other than try to break in? They should also do design analysis. As we saw above, failure does not mean the system is safe. In fact, it is unlikely that the team can test a significant number of system conditions. Penetration testing should always include design review to improve the probability of

detecting flaws in the architecture. You can reasonably ask, "If design review should always be part of penetration tests because design review is better at finding vulnerabilities, why should we do the penetration test part?" This is a valid question. I will leave the answer as a logical exercise for the readers of this paper.

What kinds of attackers will the penetration test emulate? There are, briefly, three categories of attackers: sport intruders; competitive intelligence; and foreign intelligence. Sport intruders are broken down into three subcategories: novices, apprentices, and crackers. Novices tend to stay on a single machine. They have not yet learned how to attack systems over networks. They are generally restricted to exploiting world writable files and bad passwords. Apprentices have hacker mentors. They learn to attack systems over networks, and they use tools commonly available on the Internet. Crackers write their own tools and communicate extensively with other hackers. They sometimes try to sell information they happen across. Many penetration teams will only use techniques up to low- or mid-level crackers. Detecting vulnerabilities to more advanced attacks requires a methodology that is not well understood in the penetration testing community. Be cautious if your penetration testing team claims to emulate "all" hacker skill levels; it is very unlikely indeed.

What tool sets will the team use? There are a number of commonly available tools that penetration teams will use. The Internet Security Scanner (ISS) is probably the most popular. It looks for common vulnerabilities in networked systems. Other common tools are "SATAN": another vulnerability scanner. "Toneloc": a wardialer. It looks for modems. "Modem Scan" - a 21st Century wardialer. "Strobe": a port scanner. It sequentially attempts connections to TCP ports. "nmap": a bulk port scanner. Available in Phrack 57.

Some teams use home-grown tool kits. Teams that build their own tool kits (Network Security Associates fall into this category), rather than relying on commercial products like ISS, seem to have a greater understanding of the underlying issues. After all, ISS is built and marketed for organizations to use themselves; why pay consultants to come in and use it for you? The penetration team's testing methodology document or lack thereof, is also enlightening. What are the suggested end conditions for the test? The test should have a stop state. Or, alternatively, an agreed-upon success criteria. We cannot merely say, "Stop when you get in." You

need to set specific goals for the penetration test. For example, write a file in a particular directory, elicit a response outside the design responses possible, get an interactive shell, steal some pre-determined data or probe an internal system without being detected by a firewall. The authors' chosen concluding move is, on an NT4 domain, to steal the SAM and leave a jpeg of his youngest daughter on the root of the primary domain controller. However, failure to properly define end conditions can result in unmet expectations, hard feelings, or, probably the worst of possible outcomes, a false sense of security.

Does the prospective penetration testing organisation use hackers, convicted or otherwise? A number of penetration testing teams either are run by or employ hackers. Some of the organisations use hackers that were never arrested, or were arrested as minors so that their official record is clean. Resist the temptation to sup with the enemy. Hiring hackers is an insult to legitimate security professionals everywhere, and it degrades public confidence in the profession's integrity. There is nothing a hacker knows that a well-trained security engineer will not know; you will not gain anything from hiring them, provided the rest of the team is competent. As businesses operating in the Internet community, we all have the social responsibility to raise the common standard of acceptable behaviour. Hacking is not acceptable. Rewarding hackers for their criminal activities or skills will only encourage others to take up hacking in hopes of becoming a highly-paid consultant. Security professionals should not condone or support this kind of behaviour.

Hackers are a close-knit group. They swap techniques, tools, and information about target systems. Many hackers working in "legitimate" organizations keep in contact with the underground to maintain their skills and toolsets. Some of them even continue hacking both sides of the fence. The information that hackers obtain in the course of penetration testing is very likely to leak back into the hacking community. Unfortunately, at least one of the U.K. Blue Chip auditing firms uses hackers for their penetration testing. One of the other well-known consulting organizations uses hackers as well. There is one known case where a City of London firm encouraged the CTO of a large corporation to authorise the use of hackers, in direct contradiction to the security director's instructions.

Financial institutions in the United States are under federal obligation to avoid employing people if the individual in question has been convicted of a felony or

breach of trust that would involve more than one year in prison. Bank regulators are empowered under Title 12 of the U.S. Code to call for removal of known felons or close the institution. Similar regulations apply to the securities industry. Based on discussions with federal regulators, this prohibition applies to consultants as well. Concerned industry professionals should contact their respective legal departments. Any organization shopping for a penetration test should make absolutely certain that their penetration team does not use hackers and that the organization providing the penetration services exhibits high standards of ethical and professional behaviour.

Remember not to feel too confident if the team fails to penetrate your system. If the team fails to penetrate the system, it absolutely does not in any way mean the system is secure. It means only that the team failed to penetrate the system, and nothing else. The quality of a penetration test varies greatly with the team used. Many traditional systems integrators bill their system administrators as "security experts." The practice is misleading and potentially dangerous for the target organisation. Organisations should check the credentials of any organisation testing security of their network. Some questions to ask are: Do they hire hackers? Do they have professional certification? How long have they been active in the Information Security field? How much experience do they have with the protocols and architectures in the networks under investigation? Asking these kinds of questions, and the others mentioned in this paper, will help separate the wheat from the chaff and increase the likelihood that the test will meet your expectations.

Stephen Lewer, 1999.