Penetration Testing has been a part of information security since the early 1990's, yet it is still a very much misunderstood practice – many consider it something of a 'black art'. Many CIOs and ISOs get excited at the thought of hiring a firm to perform a penetration test, because they imagine the very act of commissioning one somehow validates the idea that they and their organization are serious about security. This notion, combined with a lack of understanding of the realities of penetration testing and misconceptions about what penetration testing entails, tends to distort expectations about the penetration testing process, means and results.

In practice, there are a number of very real, very important considerations concerning scope, risk and goals which must be carefully evaluated by any organization who wishes to commission, engage in or conduct 'penetration testing'.

Definition of Penetration Testing

Time after time, security firms are contacted to 'do a pen-test' on an organization, only to discover that the organization insists that any testing can only be done under tightly constrained conditions and with highly structured Rules of Engagement, thereby defeating the purpose of the exercise.

Therefore, it is important to establish some nomenclature. What is the definition of "Penetration Testing"? NIST SP 800-42, *Guidelines on Network Security Testing*, defines it as:

"*security testing in which evaluators attempt to circumvent the security features of a system based on their understanding of the system design and implementation…to identify methods of gaining access to a system…"*

"Penetration Test" is an oft-abused term and should be confused with "Vulnerability Assessment". The goal of a Vulnerability Assessment is to determine the level of risk and exposure an organization presents on external and internal networks, devices and hosts. In a Vulnerability Assessment, risk is highly managed and impact to production systems is taken very seriously. Any possible negative impacts are factored in as an audit risk.

The goal of a Penetration Test is to *break into stuff*. To do so, the testers must necessarily pose temporarily as bad actors, and assume a hostile attack posture, in order to properly simulate real-world attack scenarios. Truly bad actors are not constrained by client requirements, uptime issues or proper authorization. While responsible pen-testers take pains to avoid any intentional negative impact while posing as bad actors, the attack toolset and techniques necessarily become more direct, and the risk of negative impact rises. As NIST SP 800-42 goes on to say:

"*Penetration testing should be performed after careful consideration, notification, and planning...is a very labor-intensive activity and requires great expertise to minimize the risk to targeted systems…the possibility exists that systems may be damaged in the course of penetration testing and may be rendered inoperable…Although this risk is mitigated by the use of experienced penetration testers, it can never be fully eliminated.*"

Here is a general illustration of the types of security testing an organization can undertake:

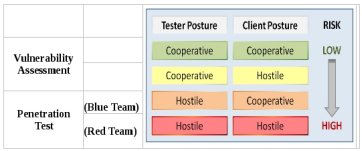

The vast majority of Vulnerability Assessments fall into the top category, Cooperative-Cooperative, with some elements of Cooperative-Hostile. These engagements usually employ industry-standard vulnerability scanning tools and data collection utilities which are largely passive and operated under controlled conditions. In these scenarios, the Rules of Engagement tend to be very well-defined, the audit risk is manageable and the overall impact is generally low.

Penetration Testing normally falls into the bottom two categories [1], where the testers assume a hostile posture and utilize a larger and more 'unfriendly' toolset, up to and including Denial-of-Service tools. Some of the techniques utilized by competent pen-testers are large-scale packet manipulation, Layer 2 protocol manipulation, buffer overflows, SQL injection, social engineering, and other techniques which are largely considered 'hacker' activities. These practices carry an element of risk which may not be suitable for certain

organizations.

Black-Box Testing

“Black Box” testing, also referred to as “zero knowledge” testing, is a scenario in which a penetration tester operates with only the barest minimum of information, about the target organization, such as a website or domain name. A common request is that penetration testers perform black box or ‘zero knowledge’ testing on a network, operating under the assumption that in order to fully simulate real-world conditions, testers should operate in the same environment and with the same constraints on information that potential attackers would. The flaw in this assumption is that the constraint on information about the target organization is usually the *only* constraint attackers operate under.

In practice, a potential attacker has virtually unlimited time to do research, reconnaissance and data-gathering on a target network. When pen-testing an organization with zero knowledge, forming a definitive picture of **all** available networks and services an organization may employ could take months of concerted effort. Additionally, due to the nature of the tools and techniques used in penetration testing, there are further legal ramifications to consider: if the testers incorrectly identify a target host or network as belonging to the client, the testers could be held legally liable for negative impact to those targets.

Additionally, attackers are not constrained by the practicality or legality of other, less technology-oriented information gathering methods such as dumpster diving, spear-phishing, pretexting, or theft. Knowing that the human element is often the weakest link in security, determined attackers employ a wide range of clandestine or social engineering attacks in order to gain information. In extreme cases, they may attempt to gain physical access to a property for the purposes of theft or network backdooring. These are common tools in an attacker’s repertoire, but options not usually available or practical for penetration testing engagements. If this type of activity is requested or required, the testers would need to research and acquire sufficient legal protections for all parties, and allocate sufficient time in which to carry out these activities, thereby incurring substantially greater cost to the client. For these reasons, black-box testing is usually not a cost-effective measure when performing a pen-test of an organization.

“Gray box” or “partial disclosure” testing is the most common type of penetration testing, and the most recommended, as it more accurately simulates a real-world scenario: the type of information which a knowledgeable attacker will be able to obtain without having to resort to time-consuming and tricky discovery methods. Usually, this is as simple as providing the testers with the exact netblocks and domain names owned by the organization. Some social engineering techniques may be employed with this type of testing, however, it is usually of a nature less likely to cause legal issues.

“White box” or “full disclosure” testing provides testers with complete knowledge of the hosts, networks, applications, ports, protocols and source code to be tested. This is the most cost-effective approach for penetration testing, as it eliminates all discovery, enumeration and footprinting requirements for the testers. There are some risks to this type of testing, however. Firstly, it does not simulate reality in any way – if

any attacker ever got this level of detail about your organization, you have bigger problems. Secondly, having full knowledge of every aspect of a system fundamentally changes the way a tester may approach attacking the system, which may run counter to the intent of the exercise: to simulate an attacker. Finally, and not least, if any penetration testing firm insists on full disclosure in order to carry out a penetration test, you should review that firm's qualifications very carefully – they may not possess the requisite skills to do a proper penetration test. White-box penetration testing is best used as a follow-up to black-/grey-box testing.

Secondary Exploitation

A central consideration of penetration testing is the depth of penetration and level of secondary exploitation desired. Some organizations may request that testers immediately cease operations as soon as verifiable compromise occurs. Others may not be satisfied if the testers compromise the entire DMZ – they require a full assessment of how deep a determined attacker can penetrate into an organization.

If full compromise and deep network penetration is desired, testers will often employ rootkits, agent-based tools and other malware on compromised hosts in an attempt to gather information and as a base from which to deploy attacks. These tools will often remain active for weeks at a time, to maximize information gathering and to test network protections and security practices. The presence of these tools increases the risk of operational impact to the compromised hosts.

'***Stealth' Attacks***

It is sometimes requested of Penetration Testers that they attempt to evade detection, or that they engage a vigilant target. There are excellent business justifications for this approach, chief of which is to determine the abilities of an organization to detect, identify and appropriately respond to intrusion attempts on the production network. However, it is not the most cost-effective approach. There are several techniques for 'flying under the radar' and they are largely time-based. Therefore, it may take days or weeks for testers to perform what would normally take hours. This can increase the cost of the effort substantially.

Successful intrusions into organizations with a properly funded, trained, prepared and alert staff are extremely improbable and exceedingly rare. The net result of such an exercise would be better viewed as a 'live-fire' training exercise or an operational evaluation opportunity rather than a true test of organizational information security – most successful full-scale intrusions occur when vigilance is low or non-existent.

Social Engineering

Another common request for pen-testing firms is that they attempt some social engineering attacks on the target organization. Social engineering comes with a high degree of risk, not only for the target organization, but for the penetration testers themselves. Securing the proper authorizations and legal waivers prior to engaging in these activities is critical, and coordination with human resources, legal departments, security organizations, and even local law enforcement may be necessary in order to adequately ensure against liability or harm.

That being said, social engineering as a part of a penetration test has some value, as it can help an organization evaluate the efficacy of security awareness training or test employee adherence to standards of conduct. It can help to reveal poor security practices, policy gaps or low security vigilance. As discussed above, the human element is often the weakest link in the security chain, and attackers take great advantage of this.

These tests of the 'human element' can be as simple as interviewing employees about security practices, or posing as a vendor to solicit information about specific hardware or software used in the enterprise. Or they can involve other, more risky activities such as: posing as an employee in order to gain passwords or remote access; 'dumpster diving' to see if sensitive data is being improperly disposed of; attempting to tailgate employees into secured areas; and even RF access card hijacking. Not all penetration testing firms will offer these services, and some of them verge on the exotic. Some of these methods cross into the realm of physical security testing. However, attackers make no such distinction. They will use any and all methods to gain access to a desired target network or system.

If requesting social engineering services as part of a penetration test, it is necessary to be very specific about the activities desired, and to adequately gauge the impact these tests may have on the organization's networks, systems and, most importantly, personnel. Employees who are successfully 'duped' by a social engineer may react negatively, and care must be taken to address the fallout from social engineering activities.

Relative Perceived Value

The concept of Relative Perceived Value is one which is not often discussed by security practitioners, but it can have a major impact on the ultimate efficacy of a penetration test. Relative Perceived Value is the difference between the perceived value to an organization of a protected asset and the perceived value of that asset to a rational attacker. Most organizations place a value on information assets which is determined by its own notion or calculation of the value of the asset. If an attacker places a substantially higher value on that asset than the organization does, then the chance that asset will be attacked rises, as does the chance the asset will be compromised by such an attack. Attackers have nearly infinite time to seek out a flaw in a system's security, and the higher the perceived value of the target, the more time and resources will be brought to bear on it. Consequentially, more effort and expense should be put into securing high-value assets and testing the protections surrounding it.

To better understand the concept of Relative Perceived Value, it is important to understand how an external attacker assigns value to targets.

As a general rule, attackers rarely know exactly what systems are in place in a given network. The fact that you just rolled out a brand new online accounting system may not register with any attackers reconnoitering your organization, except perhaps for the fact that you now have a new page on your website, or a system has just come up on one of your external IP addresses. So while the organization knows it has a new, high criticality system online, the majority of attackers will approach it as an unknown quantity with neutral value.

Since the attackers are looking at your network as a whole, they will likely probe the new site with automated attack tools, looking for obvious weaknesses. However, the overall goal is to break into your network, not necessarily subvert a particular site or host from the outside. They know that once they are past your perimeter defenses, they can begin taking a closer look at what is behind them in more detail, and reassess the value of targets visible from the compromised host. Therefore, when an attacker is evaluating an organization's external defenses, the highest value targets on the network are usually the systems that seem most vulnerable.

In an effort to reduce costs, an organization may seek to limit the scope of a penetration test to a single system or small group of systems. Because of the way that attackers generally apply value to systems, the effort may ultimately fail to give an accurate representation of the overall security of the system. The system may be externally impenetrable, but placed within reach of other, less secure systems. Attackers understand intuitively that if the front door is locked, the side window may not be. It is important to test not only the system, but its full operational context.

Relative Perceived Value is important to a penetration test in that it helps to better define the scope, thrust and intent of the effort. Penetration testers understand what attractive to attackers both from a strategic perspective (in terms of using it a 'beachhead' into a network from which they can launch further attacks) and an absolute perspective (in terms of the value of the data/function of the target). By working with the penetration testers to set the testing plan and overall agenda based on this understanding, the overall value of the effort greatly increases. Additionally, the organization can gain valuable knowledge on how to approach future security initiatives.

Negative Results

The vast majority of all successful network incursions occur due to poor configuration, known vulnerabilities which are left uncorrected, or the unwitting infection of an internal host by a user – not by some zero-day vulnerability or hacker *über*-tool. In a security-conscious organization which employs proper attention to detail, the attack surface is exceedingly narrow. Additionally, as the baseline level of security applied to devices, firmware and protocols is increased, entire classes of attacks become useless: at one time it was a trivial exercise to knock network devices offline simply by port-scanning them or sending them bad packets, in which case they would often 'fail open'. As such, there will be cases in which testers are simply unable to penetrate the network. This is not a guarantee of perfect security. Attackers are opportunistic by nature; they tend to go after the 'low-hanging fruit', and are constantly developing and employing new exploits and tools, often within hours of a vendor patch release. As new attacks develop, and as changes happen within the organizational network, however small, the security posture changes. Today's impenetrable network is tomorrow's botnet; it only takes a single vulnerability.

--

Penetration testing is not for every organization. It carries a moderate to high level of audit risk and can be expensive and time-consuming. However, done properly and with full understanding of both the risks and the benefits, it can impart great value to an organization's security posture and practices.

About the Author

Casey Priester, CISSP, CISA, CEH, SSCP is Director of Operations for Prometheus Global. He has been performing penetration testing and Vulnerability Assessments on commercial, municipal and federal networks since 2001. He can be reached at 703-266-6006 and via e-mail at casey.priester AT prometheus-group DOT com.

1 A Hostile-Cooperative effort is more colloquially known as *Blue Teaming*, and a Hostile-Hostile is known as *Red Teaming*. For more details on this as well as penetration testing techniques, see National Institute of Standards and Technology Special Publication 800-42 (NIST SP 800-42) at
http://csrc.nist.gov/publications/nistpubs/800-42/NIST-SP800-42.pdf