Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at http://www.giac.org/registration/gsec

Preparing for a Vulnerability Assessment

(Penetration Test)
By Kenneth Newton CISSP

INTRO

Penetration testing can be one of the most beneficial exercises used to validate the effectiveness of a company's information security program. While normally a stealthy process, by the time an assessment is completed, it can result into one of the most visible endeavors a security department can conduct. Visibility can be beneficial if all goes well; it can be very damaging to the development of a security program, and ultimately the security of a companies information security assets, if it doesn't.

There are many keys to conducting a successful vulnerability assessment. Most attention is normally given to the actual mechanics of an assessment. Given the risk to a company's production environment, and to the image and reputation of an information security program, it is surprising how little preparation may go into the planning of an engagement. Normally assessments are contracted out to security companies, who have experience in conducting efforts of this type. This level of experience may vary. In addition, the dynamics of an environment are so diverse, what works at one firm may be devastating at another. In the opinion of this author, the main key to the successful completion of a vulnerability assessment is in its preparation. There are so many key decisions one must make in regards to an assessment, having most of them asked and answered before starting is the best way to guarantee the least number of surprises during the event.

In the early days of penetration testing, simply justifying the need for a test was an endeavor. Since the rise of the Internet economy, and all the opportunities for exploitation that go along with it, the most basic corporate user has some comprehension of the need to test the security of an environment for common vulnerabilities. Additionally, Sept. 11[th] cast away all doubt in the need for a strong security posture.

Over the years two distinctly different, types of assessment methodologies have evolved.

PHASES: PHASE ONE

The first, most 'Hollywood', aspect of testing is best known as "Penetration Testing". This is the aspect of a Vulnerability assessment most people think about when they discuss security testing. Usually done under 'zero knowledge' conditions (where the penetration team starts with little or no previous knowledge of the environment being tested, this is the 'break in' type test. Using any means necessary, get into as many platforms as possible to show what information is

1

available to the wily hacker. This type of testing uses the philosophy of exposing the weakest link in the security chain to get access to a platform, even if the weakest link is on a different device, one possibly even using a different operating system.

In this scenario, a person could break into a UNIX system that has no known vulnerabilities (latest patch version, no un-necessary daemons running) by exploiting a common ID and password on a less-secure, Windows NT system. While this type of testing is important, for it considers the security of the entire enterprise as a single entity, it does not actually focus on vulnerabilities that may exist on every type of platform. In other words, if the NT vulnerability that was exploited was fixed, there still could be many UNIX vulnerabilities that go undetected and unresolved, just ones that aren't as easy to exploit. This does nothing to expand the understanding of what needs to be done to be more secure. What it does do, is increase awareness to a corporation of the importance of information security, and in helping that awareness, possibly allows the Information Security area more of an ability to fund its efforts in the future. In that aspect 'Penetration Testing' is a political tool that gives example to what could happen if a malicious user were to get in.

Penetration Testing is the 'James Bond' part of testing. It is what the normal InfoSec geek wants to work on, because it challenges a security auditor's technical skills and abilities, and has the highest of degree of impact and immediate satisfaction. It can also be the most revealing part of a test. For a firm to know its vulnerable is one thing, to prove to them it is, is another. Once tester is in though, and awareness to the vulnerability used to gain access is made aware, there is little knowledge of what other exposures exist on that platform. If this were the only part of a vulnerability assessment, then very little additional security may be committed to the affected platforms, for only one vulnerability was exploited. This is where testing of a diagnostic nature comes into play.

PHASES: PHASE TWO

If 'penetration testing' were described as phase one of an effort, then 'diagnostics' would be considered a good name for phase two. In this part of a test, an assessment team will work in the open, looking for many different vulnerabilities. This would include examining OS platforms, networking systems, and applications. The philosophy here is not to 'get in' but to make a thorough examination of an entity for all known weaknesses. To easily find these exposures, an assessment team may have to enlist the cooperation of a support area, or may even require escalated privileges to examine how things are secured. To use an analogy, it is much safer to see if a house will stand by looking at the blueprints rather then pushing against the side to see how much pressure it will take before it gives way.

In phase two of a vulnerability assessment, an opportunity also exists to point the assessment team towards environments not found during phase one. In this part of the test everyone has a chance to include areas they are concerned about. While it would spoil a test to inform a team during phase one, in phase two, any and all information that can root out any vulnerabilities is key to success.

External vs. Internal

Information security 101 talks about where the main threat to an organization resides. The common description defines two broad categories of threat, internal and external.

An external threat encompasses any vulnerability that penetrates the outer border of an enterprise. It also includes any servers or environments that are customer facing. In an external test, analysts would conduct a penetration test against production Internet environment. An Internet assessment focuses on application weaknesses of the storefront and customer-facing environments (business, network and customers). In fact, Graeme McLellan in his article **Alternative security penetration testing** states "Examining the security within online applications is vital, and just as key as examining the underlying network infrastructure." When one considers hackers, the first image that comes to mind is an 18 year old kid in a dark room with a PC, or somebody in Romania, or simply, the Internet. External is anything that crosses a network boundary. Points to include are modem connections, leased lines, Value Added Networks, the Internet, and any other similar connection point. In this case the world itself is the threat.

When one thinks of an internal threat they may normally think of a disgruntled employee or a malicious vendor. While that may be true, an internal threat can originate from an entity with goals of corporate espionage, or something as benign as a curious secretary browsing an improperly secured directory. The most dangerous part of an internal threat has to do with the availability of information. Inside the walls of a corporate facility one may have knowledge of systems due to their job, or be able to get information meant for employees' eyes only. This information is normally geared to provide knowledge about companies' internal practices; very helpful, but ultimately damaging if used improperly. In an effort to keep as many testers in phase one during an internal test, different locations could possibly be used for the event. The vendor could communicate with each other from these locations to share findings unless/until they are caught. If caught, the attacker(s) in that location will move on to the next phase of the test while other members in other locations could continue.

In this author's experience, most vulnerability assessments are conducted from the outside in, first external then internal, in two distinct scenarios. The first, external scenario is conducted as a person with no ability to gain information

3

about a company, due to physical location. This is the Internet/modem type connection, the storefront. The second scenario is one where a person is thought to reside in a corporate facility or satellite office or to have access to one. This scenario gives the attacker, at the least, a network connection. Sometimes, a basic logon could be provided, to simulate an employee experience.

With these two scenarios in mind, there are some newer considerations to make.

<u>Wireless LAN</u>

There are literally dozens of articles discussing the need to secure one of the newest vulnerabilities to a corporate network, wireless LANs; most notably, 802.11b. A testing process has to go in to determine a casual drive by doesn't defeat years of security implementations by including an examination of/search for unsecured wireless environments.

<u>Physical</u>

One may not consider physical aspects of Information Security when planning a vulnerability assessment. This may be a short-sighted way of looking at things, especially in this day and age. Physical protection of Intellectual Information assets could possibly turn out to be one of the most important aspects of an assessment. A hard-drive in hand is, in fact, a hard-drive owned (owned meaning 'game over' the data is yours). The best firewalls in existence can't protect you if the back door to a data center is unlocked. Since the physical vulnerability of IS assets can also be important to the security of the people occupying desks near those assets, people really understand why security exists, and why it needs to continue to progress and develop. This is an understandable situation. In a presentation, a description of how an OS was compromised will never compare to one revealing how a person was tricked into giving out information or letting someone into a building. The impact of a physical security breach is more tangible.

<u>Social Engineering</u>

Considered to be the easiest and most effective tool in the tool belt, social engineering requires nothing more than the ability to lie, and be believed while doing it. The first decision is weather to include social engineering in the test or not. If a company has a security awareness program, where employees have been briefed about the dangers of divulging information, then there is something to test. Also, without a security awareness program, allowing the use of social engineering to gain information or access about a company's environment may be too easy, not really revealing any platform vulnerabilities at all. A great deal of consideration must be made as to whether a company is ready to allow social engineering to be used. If the decision is not to use it, in an effort to have an effective vulnerability assessment, then there must be a realization that a hacker

4

wouldn't make that decision, they would proceed. Not using social engineering could itself be considered a finding of the assessment, denoting the need for a security awareness program.

There are two main aspects regarding Social Engineering experimentation.

- To gain information or access to an environment. Through general or direct methods.
  - o <u>General Social Engineering</u> may be allowed to be used at any time. Defined as reconnaissance, or information gathering, to gain a general knowledge of the company architecture. General social engineering does not include specifically attempting access to systems.
  - o <u>Direct Social Engineering</u> could be considered a tool of last resort. Any accomplishments gained as a result of direct social engineering need to be documented as such. Direct social engineering is described as the modification of rule sets or coding to gain access to equipment. Direct social engineering can also be described as access gained by forcing the release of specific information on a platform, process or application that allows access. Examples include, but are not limited to, obtaining a copy of the source code a company uses for an application, having a password changed to allow access, getting a firewall rule changed, etc.

- To test a companies employees, regardless of their role in the organization. This would be a test of a company security awareness program through use of an agreed to script. Random employees of a company are called and given the same script, with the same goals, possibly their windows password. This is important especially in a large organization, as only the IT department may be involved in a general or direct use of social engineering. Since employees may have caller ID, consider using both internal and external phone lines, to see if that matters.

<u>Home Networks</u>

A mobile corporate system, in this day and age, is no longer relegated to only connect to the company network from a desk in the building. With the advent of VPN technology and home broadband, many companies allow company assets to attach to the internet and then tunnel into the corporate network. These systems are vulnerable in certain situations under these conditions. Most vulnerability assessments don't include testing of this environment. One may want to consider a controlled test though, under phase 2 type conditions, where the home user is either simulated, or informed of the testing that will occur. This

is to ensure there are no attacks conducted on equipment or resources not involved in a companies activities.

Trophies

Gaining access to a system or an environment can be enough proof to most technical or security areas that an issue exists, or a compromise is possible. To a business area, the comprehension of this event may not take hold without proof of what may happen, of what the business impact could be. The use of 'Trophies' is an important method used to provide this needed impact. A trophy is normally evidence that a confidential business process can be accessed or manipulated by an unintended party. A screenshot of a financial interface or a confidential legal document can go very far in proving the need for information security. **Hack I.T.** states that flags can be set in an environment (as in 'Capture the Flag'). If a flag is obtained, then the information security team knows the location of the flag is compromised. It is this authors opinion that a more open method should be used, one that sets most or all points of an environment in bounds, where proof of an intrusion is actual business material. Care needs to be taken to ensure that sensitive materials are not exposed to unintended parties as a result of the test. To find a document and have its contents be known as to more than the areas responsible for it can be as or more damaging than an actual intrusion by a malicious party. Guidelines must be established before a test begins to establish how far a tester can go, and what qualifies as an acceptable trophy.

Rules Of Engagement (ROE)

The key to a successful vulnerability assessment lies in documenting an overall explanation of the event; one that also sets in place the rules that will govern it. This 'Rules of Engagement' document must be an agreed upon establishment of the practices that will be followed by both parties of the test. This need is even more paramount when considering the two phase approach. One of the major definitions needed is not just how both phases will be conducted, but when and how one will transition between phase one and two. If this is not established before the engagement begins, then the overall validity of the two phase approach may be easy to challenge, once the results are presented. A well prepared ROE, one that considers both the need for best results, and contingencies that may pop up, can be the cinch pin of a successful engagement.

The following are miscellaneous statements to consider including in a ROE document. Some may be obvious, but are still important to include, some are just suggestions. Incase there is a problem; being able to reference the ROE could potentially eliminate any argument about who made the error. If a vendor did something inappropriate, and it was in the ROE, it is obvious they made an error. In the same circumstance, if the error wasn't in the ROE, then the only thing that can be done is to ask the vendor to not do it again. While that may on some level

6

be considered the same, a violation of the ROE can effect the vendors' reputation with the customer to a much higher degree.

- There will be no Denial of Service attacks.
- Attacks may initiate from different subnets. Different IPs will be used for different types of penetration testing. *
  *This signifies to both stakeholders and the tiger team they may find attacks coming from more than one location, but not different types of attacks from one address
- Attempts will not be made to 'trip' IDS equipment (i.e. do not try to get caught). All attempts should be done in a stealthy manner in order to replicate common hacker behavior. Attackers will be told there is the potential for IDS in any environment.
- Attackers will be allowed to use Social Engineering to compromise equipment.
- The company will provide each test team member with a 'get out of jail free' card should they be caught attempting unauthorized access during the test. Once caught, the assessment team must move on to the next phase of the test without sharing information.
- The vendor assessment team must have a complete plan of attack prior to conducting the assessment (i.e. plans should not be developed on the fly).
- Vendors will be able to use any equipment they have available to them.

The Company Team

How much involvement does the Information Security department have with the assessment team during an engagement. There are different thoughts on this aspect of the test. It is easy to establish that the more involvement a company has with a tiger team during an assessment, the less chance of an error occurring. It is also easy to establish that more involvement, coupled with an ROE, could also create an overall awareness to the employees of a company that an assessment is underway.

One very large benefit to being involved in the vendor proccess during an assessment has to do with the time crunch of an engagement such as this. In a true hacker event, there are no time limits to how long one might take to break into a company. During an engagement, a vendor is contracted for a finite time. If the vendor is associated with a team of persons from the contracing company, they can observe events as they go on. While in phase one, even if the contracting group is not allowed to give up information about their company, with respect to time utilization, they may be able to ensure time is not wasted. If the vendor is seen by company representatives to be going down a 'dead end',

7

pursuing an item they know is of no benefit, they can tell the vendor to move on. Likewise, if the vendor is stopping short of pursuing an item theis is of high value, they can say, proceed. This can be very benificial.

Lastly, having an employee team work with the vendor during an assessment can allow the employees document events as they occur. This is benificial in tracking what goes on, and later as an education tool. The employee team can observe the vendor to ensure they are in compliance with the ROE. As always, it is never bad to document, document, document.

Stakeholders

As stated earlier, an ROE document is only binding if representatives from both parties agree on its contents. The decision to identify and notify key areas of an ROE for approval and suggestions may seem counter-intuitive. In fact, many tests are done without informing anyone, aside from the CIO/CSO and some key security staff. The book **Hack I.T.** describes this method as "Announced" vs. "Unannounced" testing. In the opinion of this author, there is too much at risk (time, money, assets, reputation, etc) to continue in this method. Vulnerability assessments can be conducted in a secretive manner, even if a ROE has been written and shared. An ROE document doesn't state when a test will be or what techniques will be used, it only discusses what the testing conditions will consist of and how contingencies will be dealt with.

Giving the test audience a chance to give input into the planning of an assessment is a relationship building experience. The purpose of a vulnerability assessment is to reveal what's wrong with a security infrastructure. Cooperation with operational areas, partnering, allows this to occur without the 'in your face' adversity a test can bring. No one likes to look bad in a test. But, if one knows how a test will be conducted and that they had input into how it was planned, they can consider themselves part of the solution, not part of the problem.

Contacts should be made with both the business and technical areas of an organization when stakeholders are selected. These individuals should be selected from a management and operations role, to provide balanced input from their areas. These reviewers would be the same individuals a tester would contact under certain circumstances.

If a vulnerability were discovered; one that is critical to the well being of the environment; one that would cause severe negative impact if it were exploited, a stakeholder from an effected area may need to be contacted immediately. To hold back a severe finding could be very damaging, in more ways than one. If it were to be exploited by a malicious entity before the test was concluded, the people who were informed would be looked upon as the ones at fault for not escalating it. This would ultimately affect all findings, the motivation for dealing

8

with them, and the sense of shared responsibility from the uninformed support or business areas.

<u>Vendor Selection</u>

While a vulnerability assessment can be conducted by company personnel, regulatory considerations, separation of duty issues, or technical ability may necessitate contracting a 'Tiger Team'. There are a plethora of firms in the IS industry, all offering vulnerability assessment practices. Of them all, one can consider some primary points when evaluating a vendor suitability to be a tester for their firm.

- Reputation:

    Does a vendor risk something by causing a customer pain and embarrassment during a test? This may be why most of the major financial auditing firms have 'Tiger Teams'. They survive upon their reputation as uncompromising professionals.

- Skill:

    As there are different levels of 'hacker' (Coder, script kiddie, nubie), there are different skill levels for Tiger Teams as well. As example, a firm may be an excellent Windows shop, but not have a great deal of experience in UNIX. Companies with strong networking experience may not be as experienced in web development. As such, the firm may not have the best skill in exposing weaknesses in web applications. As mentioned in the article **Net security testing requires more than tools** by Phillip Whitmore, "Tools are just that, tools. A hammer doesn't build a house, the builder does." One must pick a builder that is skilled or the house could fall.

- Experience:

    What is the experience level of a firm? A newer firm may be talented, but inexperienced. Even experienced people who haven't worked together could create an adverse experience for a customer. Much of what a customer pays for is vendor ability to take their experiences with other firms and apply them to their engagement. It is better to have done it than to have read about it being done.

For larger endeavors, it may be possible to consider more than one firm for different aspects of a test, internal and external as an example, this could be a challenge to coordinate though. A company will need to evaluate its capabilities before considering picking the most ideal firm for different parts of a test.

9

Conclusion

While the mechanics of a vulnerability assessment are often the most documented and discussed part of the assessment methodology, preparation for a test is the key for a successful test.

A vulnerability assessment can be a very exciting event. There is enough excitement to go around during an event to have to add excitement when something unforeseen happens both parties didn't consider before proceeding. Adequate preparation for an event can provide more results for the investment given and provide the ultimate excitement, a successful event that will allow an information security department to continue its program with the support of its management and employees…

Bibliography

Klevinsky , T.J.; Laliberte, Scott. **Hack I.T.: Security Through Penetration Testing** Reading: Addison Wesley Publishers 2002

McLellan, Graeme. **Alternative security penetration testing** Price Waterhouse Coopers
http://www.pwcglobal.com/extweb/ncinthenews.nsf/DocID/163C0A87887D6936CA256D03000BFF77

Whitmore, Philip. **Net security testing requires more than tools** PWC
http://www.pwcglobal.com/extweb/ncinthenews.nsf/DocID/E61258581EF399B8CA256B09007E17F4

Shipley, Greg. **Anatomy of a Network Intrusion** Network Computing Magazine, October 18, 1999
http://www.nwc.com/1021/1021ws1.html

Libicki, Martin. **Ghosts in the Machines?** http://usinfo.state.gov/journals/itps/1198/ijpe/pj48libi.htm

Patterson, Patrick **Preparing for a Security Audit**
http://www.carillonis.com/en/publications/Preparing%20for%20a%20Security%20Audit.pdf

Burrows, Dave **Penetration 101 – Introduction to becoming a Penetration Tester**
http://www.sans.org/rr/paper.php?id=266

Hayes, Bill. **Conducting a Security Audit: An Introductory Overview** http://www.securityfocus.com/infocus/1697