Considerations about Red Teaming Usage in Assessing Information Assurance

ADRIAN FURTUNA, VICTOR-VALERIU PATRICIU, ION BICA
Computer Science Department
Military Technical Academy
Bd. George Cosbuc no. 81-83, Sector 5, Bucharest
ROMANIA
adif2k8@gmail.com, victorpatriciu@yahoo.com, ibica71@yahoo.com

Abstract: Red Teaming is an advanced form of assessment that models and simulates adversary actions with the overall purpose of discovering target's weaknesses and improving its defenses. Also known as ethical hacking, penetration testing or security assessment, Red Teaming of information systems offers reliable information about the effectiveness of defense mechanisms implemented. The paper presents the Red Teaming process from both perspectives: the client and the assessor, covering various aspects like: motivation, assessment types, client benefits, client risks, assessment planning, team organization, attack preparation, execution and reporting.

Key-Words: red team, red teaming, ethical hacking, penetration testing

1 Introduction

Information systems security is of utmost importance in many organizations which process confidential information, valuable data or perform critical operations.

Verification of systems security can be done in multiple ways in the form of configuration reviews, regulatory compliance checks, basic automated tests, etc. These security checks offer a certain degree of confidentiality that the system is secure but they often provide a false sense of security. The true test of a system is 'in the wild', when it faces real attacks from motivated hackers.

Seeing the security of a system only from the defensive point of view is not the best option. We cannot know how good we are protected until an attack hits our systems. Defensive security is no longer enough to protect critical systems. In order to succeed the battle against cybernetic attacks we must deeply understand our adversaries, use their techniques to find weak points in our defense and fix them before we get hit.

The concept of Red Teaming starts from the problem of understanding the adversary and his actions. If we know the way he thinks, we can anticipate his moves and we can find appropriate ways to efficiently block his attacks.

The paper discusses the Red Teaming assessment against information systems which is an advanced form of testing the systems' defenses against realistic cybernetic threats.

Because a Red Teaming test offers valuable information for decision making within an organization, the paper also details the process from the client's perspective.

The steps needed to plan and perform a Red Teaming assessment are detailed in the next paragraphs of the paper.

2 What is Red Teaming?

The term Red Team comes from American military war gaming, where the Blue Team was traditionally the United States and, during the Cold War, the Red Team was the Soviet Union. In this context, Red Teaming is defined as teams of executives 'playing' the 'enemy' to understand what the competitive context (and competitor moves) will be in some potential future [1].

While Red Teaming can be used in multiple domains like military operations , physical security, information security, control systems, corporate strategy, this paper will discuss only Red Teaming aspects related to information and computer systems. In this context, the term Red Teaming is closely related to other terms like penetration testing, ethical hacking, tiger teaming and security testing.

We can define Red Teaming as an advanced form of information security assessment that attempts to model and simulate an adversary and his actions in order to find weaknesses in a variety of information and computer systems [2].

Another definition for Red Teaming is given by Sandia National Laboratory [3]: in information assurance, the term red teaming refers to an authorized, adversary-based assessment for defensive purposes. Here

- **authorized** means that someone with legal control of the facility, system, or entity to be red teamed has agreed to the process
- **adversary-based** means that the activity is centered on what would one or more adversaries do if they were attacking the target. This implies taking into account the adversaries' knowledge, skills, commitment, resources, and culture
- **assessment** means one is making a judgment, possibly a comparison, of the state of the target with respect to actions by the adversary
- **defensive purposes** refers to the ethical approach of the assessment. This process helps persons make informed decisions about business, about security, about computer systems, about control systems.

The team itself is one important component of the assessment and is composed of individuals skilled in performing ethical hacking [4]. They employ the same tactics malicious hackers may use against information systems, but instead of damaging systems or stealing information, the findings are reported back to the organization without producing any harm to the assessed systems.

Red Teaming of information systems is different from any other forms of assessment. Its key characteristics can be defined as follows:

- The assessment goes beyond compliance checks
- Provides a credible model of a realistic threat or adversary that can be used to evaluate the systems
- It builds a comprehensive view of the target and searches for methods to affect it
- Finds vulnerabilities before hackers do and offers a chance to fix them in advance
- Anticipates adversary moves and suggests efficient defense mechanisms
- Provides a reliable basis for decision making within the assessed organization

Furthermore, Red Teaming provides the only qualitative metrics in today's system technology discipline, thus it plays an essential role [5].

Red Teaming is a process. It has a number of phases and uses various resources like the Red Team itself, assessment tools, methodologies, facilities, training programs, etc.

3. Red Teaming assessment from the client's perspective

The client requesting a Red Teaming

assessment must have legal control of the target system and he is usually authorized for taking high level decisions in the client organization (e.g. chief information officer, chief financial officer, director, etc).

We will discuss the client's perspective in a Red Teaming assessment by answering to a few common questions:

3.1 Why should an organization use a Red Teaming assessment?

There can be multiple reasons for which an organization would need such an assessment. Some examples are given below:

- The organization has implemented a new critical system and it wants to see its resistance against a realistic attacks before putting the system in production
- The organization is designing or developing a new system and it needs a third party opinion related to security issues while changes to the system are still easy to do
- The organization needs a regular verification of its systems security
- The organization wants to measure the effectiveness of its defense systems
- The organization needs to know the impact of the potential weaknesses in its cyber defense mechanisms
- A strong security assessment is required by regulatory compliance rules

In general, when the organization needs to make informed decisions related to information security aspect it should do a Red Teaming assessment.

3.2 When is the best time to use a Red Teaming assessment?

There are at least two key points during a system's lifecycle when Red Teaming should be used. The first one is during the development phase of new systems. Here the vulnerabilities are easier to fix and may not affect other systems.

Another point is in the periodic assessment phase of the system 'in production'. This helps finding vulnerabilities and observing the response of the running system against a real-life attack.

3.3 What are the benefits for the client?

Red Teaming benefits make this kind of assessment very valuable for an organization:

- The client is shown real proof that vulnerabilities exist and they can be exploited
- No harmful actions are performed against the target
- The 'attackers' will follow a pre-approved set of rules
- The attack reveals more vulnerabilities than passive (configuration review) analysis
- The client is presented with a thorough picture of the vulnerabilities and corrective measures for each of them
- The vulnerabilities can be patched before a real attacker has the chance to exploit them

3.4 What are the risks for the client?

Because the assessment is time limited, the Red Team may not cover all attack vectors against the given target. That is why Red Teaming cannot offer 100% guarantees that all vulnerabilities have been discovered. Furthermore, any change in the target system after the assessment can modify the security state reported by the test.

The client must realize that the conclusions of the assessment represent a snapshot from a given time moment of the target system.

Other than that, active testing the systems security implies interaction with the target. Although the main concern of the team during the testing is to not produce any damage, accidents may happen depending on the team's skills and special circumstances. In this case the client can immediately talk to the assessors, identify the action that caused problems and stop that action.

3.5 What type of assessment should be chosen?

There can be multiple approaches for performing a Red Teaming assessment.

Depending on the target systems, the client should choose an *external assessment* if he wants to test the security of its systems exposed to the Internet. The test will simulate attacks of an external party.

The *internal Red Teaming assessment* should be chosen when the client wants to test its defenses against internal attackers – that already have access to the internal network.

Both types of assessments could imply technical attacks, social engineering and physical facilities attacks. The client should specify if he does not wish a certain type of test.

Depending on the amount of initial information the Red Team has about the target, the assessment can fall in one of the three categories:

- *Black box* assessment – attackers do not have any initial knowledge about the target system (e.g. attacks from Internet)
- *Gray box* assessment – the Red Team must simulate the actions of malicious users who already possess some information about the target system (e.g. a disgruntled employee, client, collaborator, etc). This is the most realistic situation.
- *White box* assessment – in this case the client wants to see how secure the target is against attackers who possess almost complete knowledge about the target. (e.g. current employees, corporate espionage, etc).

All of these types of assessments offer valuable information about target's security and they can be chosen according to client's perceived threats.

3.6 Who can be the target?

The assessment can be made against:

- A network/computer infrastructure – e.g. try to gain control of the Active Directory
- An application – e.g. test the security of an e-banking web application, payment system, ERP application, SAP, etc.
- A business process – e.g. try to disrupt the billing process from the Internet
- A facility – e.g. gain physical access to one of the company's offices and connect an access point in the internal network to access it from outside

4. Red Teaming assessment from the provider's perspective

The provider is the organization performing the Red Teaming service. The provider manages the team and the assessment process. He is also in charge of the maintenance of the team by offering trainings, preparation exercises, methodologies, facilities, etc.

We have synthesized nine steps of the Red Teaming process (from the provider's perspective) that we'll describe below.

4.1 Define assessment objectives

Ideally, the client should specify the objectives for the assessment. But these objectives are not always specified clear enough and the provider must discuss with the client any unclear aspect of the assessment.

The client should choose a testing approach (*black box, white box, gray box*) depending on the initial information known by the 'attackers'.

Depending on the location of the 'attackers' the client should choose an external assessment or an internal assessment.

At the end of this phase, the provider must know what is the target, what approach should be taken, the initial information about the target (in case of gray box or white box testing), what is the timeframe of the test and what are the limitations imposed by the client. All these aspects must be clearly written in the engagement documents.

4.2 Assemble the Red Team

The Red Team is the key component of the Red Teaming process. When we say team we mean a set of two or more individuals who interact interdependently and adaptively toward a common goal or objective [6]. In our case it is the group of specialists that will conduct the actual assessment.

There are multiple approaches in creating a Red Team. One of them is to have dynamic members that will be chosen on a project basis from a pool of experts, according to the specific knowledge required by each project. The motivation for this approach is that nobody is expert in every domain and the time taken for a general specialist to become expert in a certain domain is quite high. So the advantage of this approach is a quick team setup containing experts in the required domains. The disadvantages can be a poor communication between team members and the availability of a comprehensive pool of experts.

Another approach in creating a Red Team is to have a static group of specialists that can adapt their skills to those required by the project. The advantages of this approach include good communication between team members and efficiency during the project. The disadvantage could be a slow start of some projects because of the time required to learn the details of specific required domains.

A mixed approach would be to use a set of 'core members' of the Red Team that participate in every assessment and use 'external' specialists only when there is the need of advanced knowledge in some specific domain (e.g. Exploit writing expert, psychology expert, lock picking expert, electronics expert, etc).

4.3 Reverse engineer the target

This is a preparation phase of the assessment and is especially necessary in the *black box* and *gray box* approaches because the team knows zero or little information about the target.

In this phase the team searches passively any information available about the target, matches it with the information received from the client and tries to create different views of the target [7]:

- *System view* – the technologies, devices, operating systems used by the target (e.g. the target Company uses an Active Directory infrastructure for employees workstations but it uses Unix for its core servers)
- *Functional/Logical view* – the role and functionality of each device of the target (e.g. server xxx is used both as an email server and as a file server)
- *Physical view* - the physical location of the target and its components (e.g. server xxx is located in Datacenter from city xxx but server yyy is located in developers' room in city zzz)
- *Temporal view* - e.g. working schedule of the employees from target company
- *Social view* - information regarding the people interacting/managing the target system (e.g. number of people, age, sex, social networking profiles, email addresses, etc)
- *Lifecycle view*: the phases from the life cycle of the target (e.g. phases of a business process)
- *Consequence view* – if a certain event triggers another event (e.g. unauthorized physical access into the building causes the police to show up)

At the end of this phase the team will have a 'reverse engineered' picture of the target. The chances of the attacks are higher as the picture is more complete. The result of this phase is needed as input for the next phases of the assessment.

4.4 Create and validate attack trees

With the knowledge about the target, the team is now able to create various attack scenarios that could be performed during the assessment. All these possible attacks can be grouped into *attack trees*. The term was introduced by Bruce Schneier [8] as a way to systematically categorize the different ways in which a system can be attacked. Attack trees have been adopted in the security community and have become a standard notation for the threat analysis process.

Formalized by Mauw and Oostdjik in their paper “Foundations of Attack Trees” [9], an *attack tree* is a logical tree in which the nodes represent attacks. The root node of the tree is the global goal of an attacker. Children of a node are refinements of this goal, and leafs therefore represent attacks that can no longer be refined. A refinement can be conjunctive (all children nodes must be accomplished to reach the parent node’s goal) or disjunctive (any of the children nodes can be accomplished to reach the parent node’s goal). Figure 1 shows an example of an attack tree.

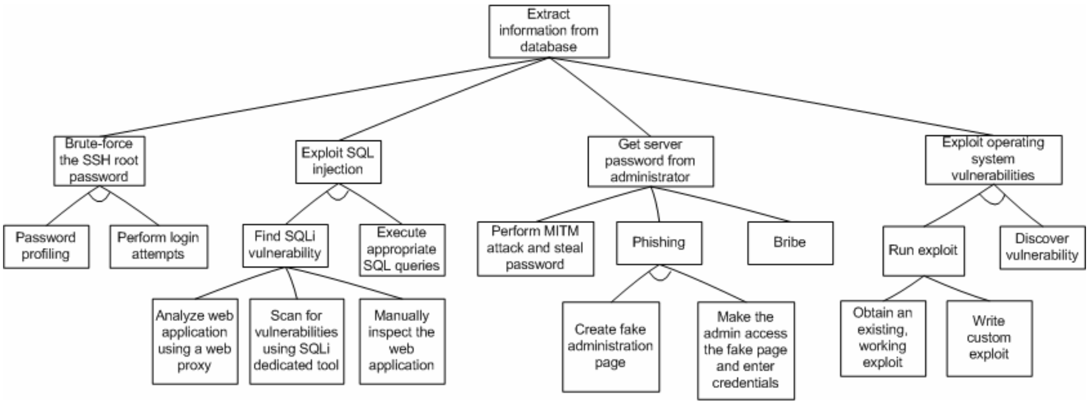
Extract information from database
Brute-force the SSH root password
Exploit SQL injection
Get server password from administrator
Exploit operating system vulnerabilities
Password profiling
Perform login attempts
Find SQLi vulnerability
Execute appropriate SQL queries
Perform MITM attack and steal password
Phishing
Bribe
Run exploit
Discover vulnerability
Analyze web application using a web proxy
Scan for vulnerabilities using SQLi dedicated tool
Manually inspect the web application
Create fake administration page
Make the admin access the fake page and enter credentials
Obtain an existing, working exploit
Write custom exploit

Figure 1 – Attack tree example

In this tree, the goal of the Red Team is to extract information from a certain database. The tree lists four possible ways to do that. This is just an example, there may be other ways to achieve this goal and the tree would get wider. Lower levels in the tree explain how these sub-goals are detailed as well.

For instance, the branch “Exploit SQL injection” requires the attacker to find an exploitable SQL injection vulnerability in a web application AND to be able to execute the appropriate SQL queries to extract the necessary data. The arc connecting the two components of this attack indicates that this is a conjunctive refinement, which means that all of the components must be accomplished in order to accomplish the SQL injection attack. Sub-attacks that do not have such a connecting arc are disjunctive, which means that accomplishing just one of them is enough to accomplish the ‘parent’ attack.

Once the possible attacks on a system have been modeled in an attack tree, the nodes of the tree must be assigned different attributes related to the attack. Bruce Schneier suggests several such attributes like (im)possibility of the sub-attack, cost, time taken, whether special tools are needed. In order to choose which attacks to perform during the test, the tree must be analyzed from bottom-up accounting each node’s attributes. For instance, the chosen attacks will be the ones for which all the nodes are marked as probable and the sum of their cost is minimum/maximum.

For each objective of the assessment the Red Team will build a separate attack tree.

4.5 Assign Red Team members to attacks

Depending on the timeframe of the assessment, the attacks can be performed sequentially or in parallel. In each case, the appropriate team members must be assigned to each attack.

For instance, the phishing or bribe attacks should be done by social engineering experts while writing custom exploits should be done by experienced persons who have already done this type of work before. The man-in-the-middle (MITM) attack and the brute-force attack can be done by a core Red Team member while the SQL injection attack should be done by a member possessing web hacking skills.

4.6 Prepare tools and methods

The time for performing the Red Teaming assessment is usually limited and it should be specified in the initial agreement with the client. So the time for the assessment is ‘expensive’ and must not be wasted with routine work. The preparation work should be done before the actual tests, during the ‘cheap’ time.

So before attempting any leaf attack from the attack tree, the team must prepare the necessary tools and test them in a dedicated environment (laboratory). Some tools need special configurations and adjustments according to the type of attack performed (e.g. update tools, configure attack options, get wordlists for brute-force attacks, download rainbow tables, etc).

For instance, the attack leaf “Create fake administration page” requires the Red Team member to install a web server and configure it so the victim can connect to it during the attack. This must be done before the actual attack.

The team can use a public methodology for testing like OSSTMM [10] or NIST SP800-115 [11], or it can use its own testing methodology.

4.7 Perform collaborative attacks

The attack tree should be executed bottom-up by the members assigned for each node.

One important aspect in the execution of the assessment is the collaboration and information sharing between team members. This is because some findings obtained by a member can be used as input for other member’s attacks. Other useful aspects of information sharing include keeping track of the project state, avoiding redundant work and finding new attack vectors.

A useful tool for effective information sharing is the Dradis framework [12]. Dradis is a self-contained web application that provides a centralized repository of information to keep track of what has been done so far, and what is still ahead. Each team member can have a user account in this application and they can work together on the same project, sharing information effectively.

The assessment ends when the objectives (root nodes) have been accomplished or when the time has expired.

The team must also document all the steps/procedures in testing in order to retrace the team’s actions in case of an incident due to testing or for retesting/verification of results if necessary.

4.8 Create the report

In the reporting phase the deliverables that will be given to the client are created. The actual work of the team during the assessment has little importance if the findings are not correctly and completely presented to the client.

The first part of the report should contain an (executive) summary which is a short presentation of the findings and risks identified during the assessment. The second part of the report must contain details about each attack performed and their results. For each vulnerability found, the report must show the associated business risks and their possible impact.

Every attack performed during the assessment must be included in the report, even if it was not successful. This can offer a certain degree of surety that the systems are safe against the attacks that did not succeed.

The last part of the report should contain corrective measures suggested for remediation of the problems found. These measures are not mandatory for the client because they often aren't made with full knowledge about client's systems and they might not fit very well in his configurations. The proposed measures should give the client a starting point for correcting the security problems in his systems.

4.9 Explain report to client

The client must know that the Red Teaming assessment offers a snapshot in time of the security state of the target.

The report contains the vulnerabilities found during the test but does not offer any guarantees that the target will be fully secure after the corrective measures included in the report will have been implemented. Any modification in client's system configuration can introduce new vulnerabilities and modify the state upon which the assessment was performed.

Anyway, the contents of the report must be personally explained to the client. This is because the decision persons / managers in client's organization often do not possess the necessary technical skills to understand the report and they may misunderstand the risks identified by the assessment.

5 Conclusions

Red Teaming is an advanced form of assessment that can be used to find vulnerabilities in a wide variety of computer and information systems. It is a process that models and simulates adversary actions with the overall purpose of discovering target's weaknesses and improving its defense.

The preparations for a Red Teaming assessment require a lot of information gathering and planning that must be done before the actual testing. After planning and execution the process continues with the report writing phase and ends with the presentation of the report in front of the client.

The assessment implies creation of different views of the target system with a thorough understanding of client's business processes.

Given the complexity of the evaluation and the highly skilled specialists of the team, the process has a great potential of finding critical vulnerabilities in the target system.

References:

[1] Beck, John C. *Responding to Global Crises Using the Change Cycle* In Thunderbird on Global Business Strategy, by Thunderbird, The American Graduate School of International Management, 384. New York: John Wiley & Sons, 2000

[2] Kraemer, Carayon, Duggan, *Red Team Performance for Improved Computer Security*

[3] Sandia National Laboratory: http://idart.sandia.gov

[4] Chris Peake, *Red Teaming: The art of ethical hacking*, SANS Institute Reading Room, July 2003

[5] Wood, Saydjari, Stavridou, *A Proactive Holistic Approach to Strategic Cyber Defense*, SRI International Cyber Defense Research Center

[6] Cannon-Bowers &Salas 1998, *Team Performance and Training in Complex Environments: recent findings from applied research*, Current Directions in Psychological Science, 7(3), 83-87.

[7] Wood, Duggan, *Red Teaming of Advanced Information Assurance Concepts*

[8] Schneier, Bruce, *Attack Trees: Modeling Security Threats*, Dr. Dobbs' Journal: Dec 1999

[9] S Mauw, M Oostdijk, *Foundations of Attack Trees*, Lecture Notes in Computer Science, 2006 – Springer

[10] Pete Herzog, *Open Source Security Testing Methodology Manual*, http://www.isecom.org/osstmm/

[11] NIST SP800-115, *Technical Guide to Information Security Testing and Assessing*, http://csrc.nist.gov/publications/nistpubs/800-115/SP800-115.pdf

[12] The Dradis Project, http://dradisframework.org/