Global Information Assurance Certification Paper

Paul Schmelzel
GSEC Version 1.3
Nessus: Vulnerability Scanning and Beyond

Introduction

In Greek mythology, Nessus was a "centaur slain by Hercules for trying to carry away Hercules' wife but avenged by means of a poisoned garment that causes Hercules to die in torment". That is the definition given by Merriam Webster's Collegiate Dictionary. Today it is more likely known as one of the best and most used security scanners in the world. Choosing your security tools is important and it is up to you to find the tools that fit your purpose. Used by hackers, system administrators, and information security engineers, Nessus is a great all-around scanner and in the following pages you will find out how to install and use Nessus. More importantly there will be an informational section on helping to understand and analyze the results that Nessus produces. After all, the analyzing is the real challenge of all scanners.

What is Nessus?

Nessus is a network security scanner. It utilizes plug-ins, which are separate files, to handle the vulnerability checks. This makes it easy to install plug-ins and to see which plug-ins are installed to make sure that your are current. Nessus uses a server-client architecture. The main server will need to be built on a supported Unix-like operating system. The client is available for Unix, Linux, and Windows. The server is not an option because "it performs the security checks" (Deraison, http://www.nessus.org/download.html). The administrator of the server sets up user accounts for other team members and issues rights to those accounts. The clients must log in to the server to be able to run their scans.

Why choose Nessus?

One of the most attractive features of Nessus in today's economy is that it is free. Anyone may download and use it, but of course the programmer is always willing to accept donations. Secondly, it is open source and many people contribute to Nessus everyday and that helps to keep it up-to-date. There will be plug-ins for new vulnerabilities within days of the vulnerabilities being released to the public. Another feature is that Nessus scans for vulnerabilities on Windows and Unix systems. This helps make it a good all-around tool so that you can scan a mixed environment in one session. Next, Nessus utilizes Nmap for port scanning. Nmap has become a standard in the Security Industry for good reasons. Nmap is known as an extremely fast port scanner. It is also very powerful, offering much more than just port scanning. Detailing everything about Nmap is beyond the scope of this paper, but if you want to learn more, which is recommended, visit the Web site at www.insecure.org. The server-client architecture is a plus if there will be more then one person utilizing the system. User rights can be defined

to lock down the types of scans they can do. If users are more familiar and comfortable with Windows, then they can use the Windows client to run scans. The plug-in architecture makes it so that each vulnerability check is an individual plug-in. This makes it easy for you to write your own plug-in. Many more features of Nessus are listed at http://www.nessus.org/features.html.
In January 2001, Network Computing ran a review of some of the top security scanners. These included NetRecon, ISS, eEye Digital's Retina, Cybercop, Sara, Nessus, and a few others. In their tests, "Nessus Security Scanner still got the highest overall score simply because it did more things right than the other products" (Forristal and Shipley, http://www.networkcomputing.com/1201/1201f1b2.html). The book Maximum Linux Security calls Nessus an "extremely versatile and up-to-date free scanner". In May/June of 2000, a survey was sent to 1200 Nmap users from the Nmap-hackers mailing list to determine their top five favorite security tools. Nessus was the most popular security tool listed. The list can be viewed at http://www.insecure.org/tools.html.

Downloading and Installing Nessus

This section will describe a typical download and install on RedHat 7.1 with X windows installed and running and Nessus v 1.1.13. This is not an endorsement; these instructions should work on most Unix-like systems. Throughout the rest of this paper a general understanding of Linux will be assumed. This setup is an example of installing the server and the client on the same machine. The obvious first step is to visit www.nessus.org and download the program. The download page is http://www.nessus.org/download.html.
You must at least download one version off the POSIX download page (http://www.nessus.org/experimental.html) because the server will have to run on POSIX system (Solaris, BSD, Linux, and others). You can choose if you want to download a stable or experimental release. It is your choice of what makes you comfortable. The author of Nessus recommends the experimental releases and that is where the link above leads. Before you download the Nessus program read the top of the web page (http://www.nessus.org/posix.html) and make sure that your system has the GTK toolkit if you plan on using the Nessus client on a POSIX system. It is also recommended that you download the latest version of Nmap because it will most likely be a newer version than what comes with Nessus. OpenSSL is recommended for a secure line of communication between the server and clients. Follow the instructions from the links to understand the other components, as this paper focuses on the Nessus program itself. At the bottom of the POSIX download page is a link to different places to download from. Pick one and make sure to download four files: libnasl, nessus-core, nessus-libraries, and nessus-plugins.
You must install the packages in the correct order because packages rely on the previous package for its install. There is a nessus-installer program but it is not recommended. It is more secure for the packages to be compiled and installed manually. This gives a little more control over what is being installed on your machine. The first package to install is the nessus-libraries. To install this package, first go to the directory where you downloaded it. Then run the following commands:

tar -zxvf nessus-libraries.x.tar.gz where *the version number replaces x*
cd nessus-libraries
./configure
make
su to root and execute: *make install*
Next you need to install the libnasl package, so make sure you are in the directory that it was downloaded to and then issue the following commands:
tar –zxvf libnasl.x.tar.gz
cd libnasl
./configure
make
su to root and execute: *make install*
Make sure to edit the /etc/ld.so.conf file by adding the line */usr/local/lib* to the end of the file. After the file is edited, run /sbin/ldconfig as root. This will make it so Nessus can find libnasl.
Now install the nessus-core package. Go to the directory that you originally downloaded the package to and use the following commands:
tar –zxvf nessus-core.x.tar.gz
cd nessus-core
./configure
make
su to root and execute: *make install*
You are ready for the final package to install now. Change to the directory that you originally downloaded the nessus-plugins package and issue the following commands:
tar –zxvf nessus-plugins.x.tar.gz
cd nessus-plugins
./configure
make
su to root and execute: *make install*
You have now installed all the components that Nessus needs to run. If you run into problems during the install, the Nessus FAQ is a great place to find answers: http://www.nessus.org/doc/faq.html.

Setting up Nessus

Now with the program installed, it must be set up for use. First, an account must be created. Use the nessus-adduser script. At the command line, run *nessus-adduser* as root. It will prompt you for a username, authentication, and rules. It will first ask for a login name and this is just what it sounds like, a user account name. The next question will ask for the type of authentication: password or certificate. Password is a good choice because you will be able to login from different machines without your certificate. Be sure to choose a good password, you do not want unauthorized clients logging into your server and using it for scans. The rules section is where you can define what scans the user can run, leave it blank to allow full access to the scanner.

Next, a certificate needs to be made for the server to work with SSL. As root, run the command *nessus-mkcert*. This script will run you through a series of questions to create your own certificate. You may answer the questions truthfully if you like or you may use the defaults.
There is a nessusd.conf file located in /usr/local/etc/nessus/ by default. This defines different settings and is also used as the default configuration when you run the server. It is a good idea to look through this file and become familiar with items you can change. For the purpose of this paper, it has been left as default.
Now it is time to make sure that all the plug-ins are up-to-date. When Nessus was installed, a script to make this easy and painless was also installed. At a command prompt run the command *nessus-update-plugins* as root. This will check the Nessus site for the newest plug-ins and install them. If you want to make sure that the most current plug-ins are installed, an easy way to do that is to visit the scripts site http://www.nessus.org/scripts.html and look at the first one on the list. Next make sure that the locate database on your system is current and run a locate for the file at the top of the list. For the example system in this paper, the commands are to run *slocate –u* first as root and then run *slocate <script name>*. The scripts may also be downloaded manually and then placed in the folder where the scripts are stored. By default this location is /usr/local/lib/nessus/plugins/. The recommended practice is to run the update script.

Running Nessus

Now it is time to start running Nessus and see what it can do. At a command prompt, as root run the command *nessusd –D***.** This command will start the server portion of Nessus and the *–D* switch will run it in daemon mode so that it runs in the background. It is recommended to run through the man pages for nessusd to understand all the options that are allowed.
Next you want to run the client piece of the software. At a command prompt, run the command *nessus*. This starts the GUI front end for the client. The first thing you must do is type in your username and password that was created in the previous step. After you log in the software is at your disposal. It is imperative that you look through all of the options and understand what the program is going to do. Also, never run a scan without permission, even if it is your own system. Your boss might not be pleased if you crash your own system running a scanner. This paper will handle the basic options that are a good starting point to learning the system.
The first tab you want to be under is "Plugins". This is where you choose what security checks you want to run. The safe step here is to click on the button labeled: "Enable all but dangerous plugins". The top part of the plug-ins list the different types of plug-ins. Highlighting one of these will populate the lower portion with each individual plug-in that can be run. You will notice that some boxes are unchecked and that they have a warning sign with an exclamation point next to their box. This is because those are considered dangerous plug-ins that could lead to system crashes. This does not mean that the ones that are left unchecked are guaranteed not to crash a system; it is just that the ones that are labeled are much more likely. Remember it is always possible to crash a

system; systems can go down from simple port scans. Double clicking on any of the individual plug-ins will bring up an information screen of what is shown if that particular plug-in is successful. You can also set the timeout for the plug-in with this window by clicking on the "Set plugin timeout" box at the bottom. It is a good idea to read through them to get an understanding of what will be checked for during a scan.
The next tab is the Preferences tab ("Prefs" for short). One thing you can do here is to set some options you want Nmap to run against your target. There are many options and this section will come from personal preference and will be a good starting point in scanning. You can choose the type of scan you would like to perform. The two most popular are the connect scan and SYN scan. By default, Nessus uses connect scan. SYN scan is also a good choice because in past experiences it has performed faster than the other options. However, some users have experienced problems when using the SYN scan. Try different options and find what works best for your situation. "Identify the remote OS" is an excellent option and a very powerful feature of Nmap. Nmap "uses TCP/IP fingerprinting for remote OS detection" (Fyodor, http://www.insecure.org/nmap/nmap_documentation.html). You can set a specific source port for your scans and this is useful if you think that you can get through a firewall with a certain port as your source. An example would be that many firewall configurations allow DNS traffic through on port 53. There are many options unrelated to Nmap further down this tab and it is a good idea to review them. For the purpose of this paper, the rest of the options are left as default.
The next tab is the "Scan Options". Here is the place where you can define some performance options. You can choose the port range to scan and the number of machines to be tested at one time and the number of checks to perform at the same time. For the ports range, it is a good idea to check the entire range of 1-65535 ports. This will take more time, however many backdoors can run on the port of a hackers choice and hackers know what ports administrators look for to be open, so the hacker could have a backdoor on a very high numbered, rarely used port. By scanning all of the ports, you will be able to know if there is one opened that shouldn't be and then it can be investigated. The default selections for the number of hosts and the number of tests to run at the same time should be fine. They may be altered if you feel your scanning machine or the machine you are scanning or your network would be unable to safely handle these numbers.
The next window is one that can be very dangerous. This is the "Target Selection" tab. Here you type the IP address(es) of the machine(s) you want to scan. It is extremely important that you type in the correct IP address and double check it before you start the scan. You may find yourself in legal trouble if you scan someone else's machine. You can enter as many IP addresses as you want or you can have a list of them in a file somewhere and just point Nessus to it. If you enter multiple IP addresses in the top box, make sure to separate each host with a comma. When you are finished and are positive that everything is right and have been given permission to run the scan, click the "Start the Scan" button at the bottom of the screen. A new screen will pop-up that will be the scan progress screen. There will be a picture of a computer with the IP address being scanned beneath it. Next to it will be two bars: the top one is the progress of Nmap and the lower bar is the progress of the vulnerability checks. Beneath the bars will be text

telling you specifically what Nessus is running at the moment. There is also a "cancel" button on the right side for you to cancel individual scans or you can select the "stop all" button on the bottom of the screen. The scan will run and when it is finished, the report will be brought up on the screen. You should not leave a scan unattended just in case something goes wrong; you need to be there to shut the scanner down.
The rest of tabs are not used as often. The next tab is the "User" tab. This will list the rules for the particular user and allow rules to be added by clicking the "Add rule" button on the right. The tab after that is the "KB" tab which stands for knowledge base. Using this option will save information that was collected on hosts from scanning. This information can be used later for further scanning. It is best to always fully scan a host again as if you have never scanned them in the past. Using the KB option could allow Nessus to re-use a port scan from the same host that was scanned early. This is bad because there might have been a backdoor installed recently and without the port scan, it would never be found. The last tab is the "Credits" tab where you can see the author and find the version of Nessus you have.

Looking at the results

The results window will be broken into five different sections. The top left section will contain the subnets that were scanned. Selecting any of the subnets will populate the lower left panel with the individual IP addresses that were scanned. Selecting an IP address will fill the top middle panel with all the ports that were discovered open on the target machine. Ports may have one of the following icons next to it: a red circle with a hyphen in the center, a warning icon with an exclamation point, or a light bulb icon and each one corresponds to a severity level ranging from highest to lowest, respectively. The red circle means that there is a serious security hole associated with the port. The warning icon means a security warning was discovered but is not as serious as the red circle. The light bulb means that is a security note on that port and this is the least serious of the warnings. If there is no icon next to the port then Nessus found no security issues with the port. This is not a guarantee that there is nothing wrong with the machine. It is important to remember that Nessus can only check what it has been told to check for. Selecting a port will fill the upper right window with the warnings discovered on that port. Selecting one of those will fill the bottom right of the screen with the description of what was found. It could be information the machine leaked or a vulnerability that was discovered. Many times a solution will be presented at the same time.
At the bottom of the screen, there is a button to save the report. Clicking on this button will bring up a save screen. You can name the file, decide where to store it, and choose the format for it to be saved. The report should first be saved in the native format for Nessus that is .NBE or .NSR is older versions. This is a good way to save it because that way you will always be able to load the report and view it in the same manner that you just have. Then it is a good idea to save it again with the same filename, except save it in html format. This is nice to have to send to other people that might need to view the results of the scans and most users will have access to a web browser. Users will not need Nessus installed on their machines to view an html report. Being in html format

also allows you the ability to edit the file and you can remove false positives or information that you do not want to have show up. You can also add items to the file and put in your own personal notes.

Analyzing the results

Now comes the part that no one can teach you and that only research, practice, and understanding can make you better at. So far everything has been pretty straightforward and most people would be able to do it. The next part is where you will set yourself apart from others. This is where you have to analyze the results. Every scanner that has been made so far will produce false positives. False positives are alerts in the report showing a vulnerability when the vulnerability is not really there. Analyzing the results will take time but you will become more efficient as you use Nessus and get used to seeing the alerts. No one is able to teach a person how to analyze the results, this paper will simply give some general guidelines of ways to test some of vulnerabilities you will come across and point you in some directions to find out about others. The first step to analyzing the results starts before you even download Nessus. It is one thing to keep Nessus current by installing the latest plug-ins and it is another thing completely to keep yourself current. As often as possible, check security sites and try to see what is going on in the world around you. There are new vulnerabilities discovered almost everyday. If you read about a new vulnerability one day and see it come up in one of your scans next week, you will have a much better understanding and will be better prepared to deal with it. It will be very time consuming if you have to research every vulnerability that you see in your scans. One good site to add to your favorites is of course www.sans.org but there are also many others. Other sites include: www.hackinthebox.org, www.securitynewsportal.org, www.cert.org, www.securityfocus.com, and www.securiteam.com. At the SecurityFocus site you can also link up to the Bugtraq (http://online.securityfocus.com/archive/1) mailing list. This is a mailing list that subscribers post to with vulnerabilities discovered and responses to each. It is also a good idea to subscribe to some of the mailing lists, but be careful because many of them are high volume lists. Each site will be better in certain areas but viewing these sites will keep you in the know and give you wide range of information.

The next step is one of the easier checks to make for vulnerabilities that Nessus finds. These are warnings that are related to certain versions of software. When you disable the dangerous plug-ins, Nessus will not attempt to exploit certain vulnerabilities because they could lead to a denial of service. In these cases, Nessus might report that you could be vulnerable because you are running a certain version of a piece of software. It is very easy if it is your system to check the version of the software that you are running and see if you are truly vulnerable. If you are vulnerable then you can patch it if you deem it necessary and if you are not vulnerable you can leave that out of your report because it is a false positive. You could also leave it in and show your superiors that you were ahead of the game and have already patched your system. If you are not the administrator of the system, then you should contact the individual that is the administrator and talk to them about which version they are running.

The next check that can be run is to verify all the ports that Nessus reported as being open. Telnet is an extremely helpful tool when it comes to checking for ports. At a command prompt just type *telnet <ip> <port>*. This will allow you to check what banner you are presented with upon connection. This banner alone can be a vulnerability if gives away information about the program that is listening on that port. The banner might include the program name and version that is listening, which is a good start for a hacker to start researching vulnerabilities in your system. When you telnet to a port you might need to hit enter a couple of times to get a response from the system. Banners are not always a guarantee that the correct information is being displayed. It is common practice to change the banners to display the wrong information to trick a hacker into going after the wrong application. The banner could give false information about the version number or even about the program that is listening on the port. One example of using telnet to test a frequent vulnerability is to verify that SMTP relaying is possible. This is something that should definitely be checked out because if SMTP relaying is possible then spammers can use you to bounce their e-mails. Spammers will send e-mails through your SMTP server using a fake from address and send it out to a lot of people. The receivers of the e-mails will be able to see that they came from you SMTP server and this can lead to your site being black listed. Being black listed means that many places will no longer accept mail from your server, which can be a black eye to your company. A good check is to start by getting a free e-mail account from a site like Hotmail or Yahoo. Then telnet to your SMTP server with a command like *telnet <ip> 25* with 25 being the port that SMTP listens on. When you connect you will get a banner and now just start typing the following commands:
mail from: test@testing.com
rcpt to: your outside e-mail account
data: testing
.
<ctrl> d
If the server gives you a warning after the first command is entered, you can try to run the following command first:
helo <domain>
If it stops you anywhere else then you are safe from relaying. However if it says "message sent" then you need to keep checking your outside account for that e-mail. If it never comes, it probably means that you are safe because some SMTP servers will accept the previous commands and then just drop them when it checks it's rule set. However, if you do receive the e-mail then you are open to spammers and you need to follow your vendor instructions on how to close it down. Telnet has many other uses and this was to show you one example and hopefully this will lead you to find other paths on your own. Another useful check is for vulnerabilities that you have heard of but do not know how to test if they are exploitable on your system. The best way to test this is to look for programs that have been written to exploit these vulnerabilities. This can be a very dangerous method and might be one to skip if it makes you nervous. Downloading code and exploit programs can be dangerous because they could contain malicious code of their own to compromise your system or they could damage the system you use them

against. This technique is to be used at your own risk. If you do download code it is best to stick with sites that you are comfortable with and that are popular. Most of these sites do not get to check all of the code that is submitted to them but if there are programs downloaded from popular sites with a bug in them, it is much more likely to be discovered and publicized. One site: www.securityfocus.com keeps a list of a lot of useful tools. Go to their site and at the top of the page you can use keyword searches and then select "tools" or "vulns" from the drop down menu. If you cannot find what you are looking for there, another site: www.securiteam.com offers an array of different exploit code for many vulnerabilities. At their home page you can use the keyword search and then select "tools" or "exploits" from the drop down list. It is important to remember that the Internet is a great way to find items but you must also be careful in who you trust. The last check is when you get an alert from Nessus that you do not have any idea what it is or how to test for it. This is when the Internet search engines will become your best friend. When you know nothing about a warning, the best thing to do is search for it on the Internet and see what is out there. Ninety-nine percent of the time, you will find your answer there. Two search engines that have been very helpful are www.google.com and www.dogpile.com. There are also many sites that keep growing databases of vulnerabilities as they are publicly announced. Many other sites keep extensive databases of vulnerabilities. A few examples of these sites are www.cert.org, www.securityfocus.com and www.securiteam.com. The last site to mention is the ICAT metabase at http://icat.nist.gov/icat.cfm. ICAT is a searchable index of common vulnerability and exposures (CVE).

Conclusion

Running a security scanner against your systems is a very important part of the job. It is a system administrator or security officer's job to keep their systems secure and the data contained in them safe. Hackers have access to all the same information and tools that the rest of us do. Hackers run the very same tools and it is advantageous to know what the results are that they see if they scan your system. They find time to do the research, so we must also. Nessus provides a lot of functionality in one tool. It utilizes Nmap, easy to update plug-ins, and nice reporting tools for upper management. It is has repeatedly scored high on comparisons between scanners including commercial scanners that come with a hefty price tag. And of course as budgets tighten, remember Nessus is a free tool. The only cost is the users time in learning it and using it, but that is a cost associated with all tools. And luckily Nessus is an easy to learn tool. Using this tool and seeing the vulnerabilities will help you gain knowledge of your systems and help teach you how to protect them. Only by scanning your systems consistently can you keep them secure. It is important to remember that new vulnerabilities are released almost daily. All of the practice will make for more efficient analyzing skills to find out what is real and important. Finding the vulnerabilities before hackers do is a great first step in keeping you systems and company data safe.

References

1. "Best of Linux: Nessus." Linux DaveCentral. 2001, http://linux.davecentral.com/articles/view/1120/
2. Deraison, Renaud. "First Step: Install Nessus." The Nessus Project. 1999, http://www.nessus.org/demo/first.html
3. Deraison, Renaud. Nessus Features. 2000, http://www.nessus.org/features.html
4. Forristal, Jeff, and Shipley, Greg. "Vulnerability Assessment Scanners." Network Computing. January 8, 2001, http://www.networkcomputing.com/1201/1201f1b1.html
5. Fyodor. "Remote OS detection via TCP/IP Stack FingerPrinting." 1998, http://www.insecure.org/nmap/nmap-fingerprinting-article.html
6. Fyodor. Nmap. 2001, http://www.insecure.org/
7. Maximum Linux Security: Second Edition. Indiana: SAMS, 2001.
8. Merriam Webster's Collegiate Dictionary: Tenth Edition. Massachusetts: Merriam-Webster, 1993.
9. van der Kooij, Hugo. Nessus F.A.Q. 2002, http://www.nessus.org/doc/faq.html