

Workbook

SANS

THE MOST TRUSTED SOURCE FOR INFORMATION SECURITY TRAINING, CERTIFICATION, AND RESEARCH | sans.org

Copyright © 2016, The SANS Institute. All rights reserved. The entire contents of this publication are the property of the SANS Institute.

PLEASE READ THE TERMS AND CONDITIONS OF THIS COURSEWARE LICENSE AGREEMENT ("CLA") CAREFULLY BEFORE USING ANY OF THE COURSEWARE ASSOCIATED WITH THE SANS COURSE. THIS IS A LEGAL AND ENFORCEABLE CONTRACT BETWEEN YOU (THE "USER") AND THE SANS INSTITUTE FOR THE COURSEWARE. YOU AGREE THAT THIS AGREEMENT IS ENFORCEABLE LIKE ANY WRITTEN NEGOTIATED AGREEMENT SIGNED BY YOU.

With the CLA, the SANS Institute hereby grants User a personal, non-exclusive license to use the Courseware subject to the terms of this agreement. Courseware includes all printed materials, including course books and lab workbooks, as well as any digital or other media, virtual machines, and/or data sets distributed by the SANS Institute to the User for use in the SANS class associated with the Courseware. User agrees that the CLA is the complete and exclusive statement of agreement between The SANS Institute and you and that this CLA supersedes any oral or written proposal, agreement or other communication relating to the subject matter of this CLA.

BY ACCEPTING THIS COURSEWARE YOU AGREE TO BE BOUND BY THE TERMS OF THIS CLA. BY ACCEPTING THIS SOFTWARE, YOU AGREE THAT ANY BREACH OF THE TERMS OF THIS CLA MAY CAUSE IRREPARABLE HARM AND SIGNIFICANT INJURY TO THE SANS INSTITUTE, AND THAT THE SANS INSTITUTE MAY ENFORCE THESE PROVISIONS BY INJUNCTION (WITHOUT THE NECESSITY OF POSTING BOND), SPECIFIC PERFORMANCE, OR OTHER EQUITABLE RELIEF.

If you do not agree, you may return the Courseware to the SANS Institute for a full refund, if applicable.

User may not copy, reproduce, re-publish, distribute, display, modify or create derivative works based upon all or any portion of the Courseware, in any medium whether printed, electronic or otherwise, for any purpose, without the express prior written consent of the SANS Institute. Additionally, User may not sell, rent, lease, trade, or otherwise transfer the Courseware in any way, shape, or form without the express written consent of the SANS Institute.

If any provision of this CLA is declared unenforceable in any jurisdiction, then such provision shall be deemed to be severable from this CLA and shall not affect the remainder thereof. An amendment or addendum to this CLA may accompany this courseware.

SANS acknowledges that any and all software and/or tools, graphics, images, tables, charts or graphs presented in this courseware are the sole property of their respective trademark/registered/copyright owners, including:

AirDrop, AirPort, AirPort Time Capsule, Apple, Apple Remote Desktop, Apple TV, App Nap, Back to My Mac, Boot Camp, Cocoa, FaceTime, FileVault, Finder, FireWire, FireWire logo, iCal, iChat, iLife, iMac, iMessage, iPad, iPad Air, iPad Mini, iPhone, iPhoto, iPod, iPod classic, iPod shuffle, iPod nano, iPod touch, iTunes, iTunes logo, iWork, Keychain, Keynote, Mac, Mac Logo, MacBook, MacBook Air, MacBook Pro, Macintosh, Mac OS, Mac Pro, Numbers, OS X, Pages, Passbook, Retina, Safari, Siri, Spaces, Spotlight, There's an app for that, Time Capsule, Time Machine, Touch ID, Xcode, Xserve, App Store, and iCloud are registered trademarks of Apple Inc.

Governing Law: This Agreement shall be governed by the laws of the State of Maryland, USA.

SANS Auditing: System and Network Auditing

Hands-On Exercises

Revision 8.1
Q3 2016

SANS Advanced Systems Audit Workbook

Preface	6
Workbook Conventions	6
Day 1	8
Exercise 1: Samples and Errors	8
Part 1: Finding a Sample Size	8
Exercise 2: Virtualization Security and Configuration	11
Pre-Lab: Preparing Your System	11
Part 1: Starting ESXi	13
Part 2: Examining the Server	15
Day 2	21
Exercise 1: Switches and Network Symptoms	21
Part 1: Wireshark – Introduction	21
Part 2: Wireshark – Layer 2 Analysis	27
Exercise 2: Network Behavior	31
Part 1: Virtual Machines	31
Exercise 3: Routers, Firewalls, and Validation	42
Part 1: Nipper	42
Part 2: Firewall Validation	48
Part 3 – Automated Firewall Validation	54
Exercise 4: Finding Wireless Clients	57
Exercise 5: DNS and SMTP	60
Part 1: DNS	60
Part 2: SMTP	65
Exercise 6: Network Discovery and Population Management	68
Day 3	77
Exercise 1: Intro to HTML and HTTP	77
Part 2: Man-in-the-Middle Proxies	81
Install JRE	81
WebScarab	82
Burp Suite	86
Exercise 2: Application Fuzzer	91
Exercise 3: Brute-Force Authentication/Credential Exposure	102
Configuring Burp	102
Exercise 4: Session Analysis	112
WebScarab	112
Burp Suite	115
Complete Audit	120
Input Validation and Manipulation	120
Hidden Content	121
Cache Prevention	122
Session Tracking	122
Web Application Auditing – Typical Homegrown	124
Day 4	127
Lab Overview	127
Special Note About Windows Utilities	128

SANS Advanced Systems Audit Workbook

Connecting to the Cloud	129
OnDemand Cloud Access	130
Live Conference Cloud Access.....	131
WMIC Exercise	132
Alternative Approach – Remote Server.....	134
Basic Scripting Exercise	135
Basic System Information/Open Ports and Running Services.....	140
Part 1 – Install System Internals’ PSTool.....	140
Part 2 – Use PSTools to Obtain Basic Information about the System.....	140
Part 3 – Identify Open Ports.....	142
Part 4 – Use netstat to Identify Services Listening on Each Port	143
LDAP/DSQuery	146
Install DSQuery	146
Running Queries	147
Using “-filter”	148
Users, Groups, and Passwords.....	151
Part 1 – Install Cain and Abel.....	151
Part 2 – Create Sample Accounts on Windows	151
Part 3 – Use Cain and Abel to Extract the Local Password Hashes	153
Part 4 – Domain Credentials and Cain.....	155
Protecting Data.....	158
Part 1 – Use Somarsoft’s DumpSec to Extract Information about Permissions and Shares.....	158
Security Configuration and Analysis	162
Part 1 – Create a Microsoft Management Console (MMC).....	163
Part 2 – View the Available Templates and Options.....	167
Part 3 – Modify a Template to Match a Specific Security Policy	167
Part 4 – Import Your Template into the Analysis Database	169
Part 5 – Analyze Your System.....	171
Part 6 – View the Analysis Results.....	172
Part 7 – OPTIONAL: Customize a Template	173
Day 5.....	177
Section 1A.....	178
Exercise 1: Exploring Unix.....	178
Exercise 1B: Basic Scripting	184
Section 2.....	190
Exercise 1: File Integrity Checking	190
Exercise 2: Baseline Network Configuration	192
Exercise 3: Startup Scripts	195
Exercise 4: Tiger	197
Section 3.....	199
Exercise 1: Unix Log Files	199
Exercise 2: Log Analysis with SWATCH	201
Exercise 3: Password Assessment	211
Section 4.....	213
Exercise 1 : Building a Tools CD	213

SANS Advanced Systems Audit Workbook

Exercise 2: A Unix Conformance Audit.....	217
Day 6.....	224
NetWars: Audit the Flag!.....	224
Appendices.....	225
Selected Answers.....	226
ESXi Issues.....	226
Firewall Validation – Optional.....	228
Automated Scanning and Analysis.....	228
WMIC Cheatsheet.....	235
Aliases.....	235
Useful DSQuery Formulae.....	238
User Account Control Bit Values.....	240
Sample Scripts.....	242
Unix Audit Script.....	242
Troubleshooting.....	246
How do I change the keyboard layout settings inside of the virtual machines?.....	246
Web Application Audit Checklist.....	247
Basic Configuration.....	247
Authentication.....	247
Session Management.....	247
Input.....	248
Output.....	248

Preface

The exercises contained in this book are designed to go hand in hand with the SANS Audit 507 course. Each of the exercises is intended to give you hands-on experience using the tools and techniques described in the course and to help you think more deeply about what you are seeing so you can perform some deeper analysis of the systems with which you interact. In order to complete the exercises in this workbook, you need a computer with VMware Player 7.0 or higher¹. We provide a redistributable version of VMware Player on the course USB and directions for its installation.

As an alternative, you might install VMware Workstation version 9.0 or higher or VMware Fusion 7.0 or higher if you are using an Apple computer. If you do not have a valid license for VMware and would prefer to use the Workstation or Fusion product for your platform, we would recommend that you obtain the demonstration version of VMware from www.vmware.com and a license key. This version is fully functional for 30 days and is more than sufficient to complete the labs.

To complete the exercises, you also need the USB that is provided with the course materials. This USB contains a number of VMware “images” (live systems or virtual machines) in addition to tools, spreadsheets, and other materials that you will use to complete the exercises.

This book is organized so that the exercises follow the courseware fairly closely. In the event that you do not have the time to complete the exercises or if you simply choose not to do them during the conference, the exercises can be performed against systems in your own environment using the tools provided on the course USB or using the VMware images on the USB. In addition, we have taken the time to record video walk-through versions of nearly all of the hands-on exercises to provide you with examples that you can work along with. The servers used in class, except for a few additional web-based applications, are default installs commonly found in the “real world.”

Workbook Conventions

In the text of the workbook, we use a few simple conventions, as listed in the following:

- ***Text that is in bold, italic represents actions, activities, or commands that you must complete to successfully complete the lab. If you are in a hurry, you can simply use the bold, italic text.***
- Text that is in a normal font represents additional information and explanatory material. Although you can complete the lab without reading this text, we recommend that you actually *do* read this text. It will assist you in reasoning out *why* you are doing certain things and help you to analyze systems on your own.

¹ Please note that VirtualBox **will not** work with some of the virtual machines in this class. Several of the machines are “linked clones,” which means that they share a common hard drive image. This is not supported by VirtualBox.

SANS Advanced Systems Audit Workbook

- Text that is in the courier font represents actual commands that should be typed exactly as they appear at the relevant command line.

We hope you enjoy the exercises and find them beneficial. Without further ado...

Day 1

Exercise 1: Samples and Errors

Time Required: Approximately 10 minutes

Purpose: Acquaint the learner with effective methods for calculating statistically valid sample sizes and demonstrate how to arithmetically determine the margin of error for an arbitrary sample size.

Part 1: Finding a Sample Size

Please begin by locating the “Calculate Samples and Errors” spreadsheet on your USB. The file can be found on the USB stick in “Days 1-5\Useful Spreadsheets\Calculate Samples and Errors.xlsx”. When you locate the file, please double-click it to open it.

- AD Schema.xls
- Calculate Samples and Errors.xlsx
- data warehouse bram.xls
- implprem.xls
- IT RISK ASSESSMENT TEMPLATE.xls
- Risk_Register_aligned_ASNZS4360.xls
- Top 20 Critical Controls - Greg Ahira.xlsx

If you do not have Microsoft Excel installed on your system, please just follow along because your instructor will walk you through this lab. We did not include Excel in the laptop requirements because we use it for only one quick lab all week long.

Let’s begin by attempting to calculate a sample size. To do this, we will use the top portion of the spreadsheet. The top portion of the spreadsheet is specifically for determining the appropriate sample size for a test.¹

Size of sample required for level of confidence:

Population size (P):

Expected Occurrence (p):

Acceptible Margin of Error (D):

$$S = \frac{z^2 (p(1-p) / D^2)}{1 + (z^2 (p(1-p) / D^2) / P)}$$

If the expected occurrence is not known from a previous analysis, always begin by assuming a value of 50%. Following each analysis, this value can be iteratively improved for a more accurate and smaller sample

Result Confidence:	Z =	n	Sample Required
90%	1.645	321.3405	298.4515519
95%	1.96	456.19	411.3986083
99%	2.575	787.3867	662.8278127

Let’s work through a few problems.

- You have been asked to determine the number of systems that must be examined to produce useful results concerning users with local administrator rights. The last time that this question was examined, 25% of systems had users in the local

¹ Some have asked what the column labeled “n” represents in the spreadsheet. In our spreadsheet, the “n” column represents the numerator of the equation pictured to the top, right of the spreadsheet.

SANS Advanced Systems Audit Workbook

administrators group, despite corporate policy to the contrary. If the organization has 25,000 Windows desktops, how many systems must we examine if we would like to determine compliance within a 3% margin of error?

To answer these first questions, you must simply fill in portions of the spreadsheet. **First, set the population size (P) to 25,000.** This number is easy to determine.

You can skip the next field, Degree of Confidence. This is actually an artifact of an older version of the spreadsheet. If you look at the lower portion of the top section of the spreadsheet, you will see that the resulting confidence data has been broken out into the three most common confidence intervals that are used.

Based on the information provided, it appears that this test was performed before. In that test, 25% of the systems had users in the local administrators group. **Enter the value 25% in the "Expected Occurrence" field.** This field requires a little bit of explanation. You may look at it and think, "Well, we really don't expect anyone to have this, do we?" Although that might be the expectation, unless we have facts to support it, we cannot assume zero. Instead, if you have absolutely no evidence, you must assume a 50% failure rate. If you remember that we are talking about standard distributions and think about the standard "Bell Curve," this makes a lot of sense because it would be the largest portion of the population. In our specific case, we do have the results of a previous assessment. During that evaluation, it was determined that 25% of systems had the configuration in question. This means that we can plug 25% into this field.

When you change this field to 25%, you will likely notice that the sample sizes jump up in size. **Please set the "Acceptable Margin of Error" field to 3% to match the requirements.** The margin of error is the wiggle room that you are willing to accept in your results. If you are paying attention when you change this value, you will likely see the sample sizes drop by more than half!

The final part of this question requires that we provide the number of systems that should be examined in our sample. To answer this, we have to explain the concept of a Confidence Interval. Here is a way that you can state the data that may help:

"I am 95% confident that if we examine only 776 systems, we can determine the number of systems that are misconfigured within a margin of error of 3%."

It truly does indicate how "confident" you are in the results. You can see that at a 99% degree of confidence, the sample size nearly doubles. By comparison, accepting a lower degree of confidence requires fewer systems in the sample. What is the right answer to the question? **Almost all examples of this type of statistical analysis will provide results at a 95% confidence interval.** Unless someone specifically asks for something else, or if you cannot handle the size of a sample, you should generally use the 95% confidence interval for all problems of this sort.

Answer: 776 systems

2. One of your co-workers has interviewed 98 users in your organization about a specific security issue. Of those 98, four were not aware of the correct policy that applied to the question being considered. ***Please estimate how extensive this problem is in your organization. You do not have access to the total number of employees.***

Answering this question requires us use of the bottom portion of the spreadsheet. In this case, we have been provided with an arbitrary sample size. We have been asked to take this sample and turn it into meaningful data. To do so, we simply fill in the blanks!

Start by putting the number 98 into the sample size field. Next, simply insert the number of individuals who either match or did not match the criteria in the “matching criteria” box.

With these two boxes filled out, the calculation can be performed. It seems extremely counter-intuitive, but there are a lot of statistical theories and proofs behind the formula that is applied. Of course, the larger the sample size, the more accurate the results. You will see this reflected in the margin of error. In addition, be aware that as the number matching the criteria approaches the mean, you will also see the margin of error grow geometrically.

In our case, the answer is easy to determine. Because we already know that we should likely be reporting at a 95% confidence interval, we can state the following: ***We are 95% confident that 4.08% of our employees do not understand this particular policy with a margin of error of 3.92%.***

Answer: Between 0.16 and 8% of employees do not understand this policy. Alternatively, 4.08% of employees with a margin of error of 3.92% do not understand this policy.

Exercise 2: Virtualization Security and Configuration

Time Required: Approximately 35 minutes

Purpose: Familiarize the learner with the management interface for VMware ESXi. Identify common security configuration issues.

Pre-Lab: Preparing Your System

Purpose: This exercise is included to provide instructions on how to configure VMware on your system so that you can successfully complete the workbook exercises that follow. The virtual system, ESXi, that we will use this afternoon is found in the “Days 1-5/507 Systems” folder. *You will find that it runs fastest if you copy the virtual machine to your laptop before attempting to open it.* Although it will certainly run from the USB drive, it will run *very* slowly.

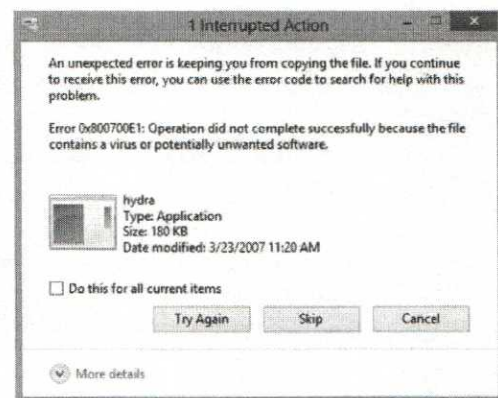
Prior to attending the conference, a reminder email was sent to the address on file with SANS reminding you about the laptop requirements for the class. These requirements indicated that you should have a computer with at least 2 gigabyte of RAM running a 64-bit version of Windows 7 Business (or higher). Your user account must also be in the Local Administrators group on your system for a number of the exercises during the course. In addition, your system should have at least 5 gigabytes of free space on the hard drive. If you have any problems or questions throughout the class regarding any of the exercises, please ask the instructor for assistance!

Many of the exercises that we will perform in this course assume that we have a virtualization platform installed along with a number of virtual machines. To prepare for those exercises, let's get the required software and the virtual machine installed. *To do this, please insert the course USB now if you have not done so already.*



IMPORTANT NOTE:

When you insert your course USB, it is very likely that your antivirus software will report that several of the files on the USB are suspect. The files most commonly reported on are PSTools and all_attack.txt. It is possible that your software will report on more or less than these. It is also possible that the antivirus software will not immediately report but instead will simply generate errors when you attempt to copy these files or that in the process of copying these files to your system, your antivirus software will delete them silently. Please note that these tools will not harm your system in this course. If you want, you can temporarily disable your antivirus software, but please do not violate any of your corporate policies in the process. These tools will have little or no impact on your experience in this course.



If you do not yet have VMWare Player installed

With the course USB for days 1-5 inserted, please locate the VMware Player installation executable¹ in the root directory of the USB. Please double-click this file to begin the installation process.



When the installer is running, VMware will walk you through an installation wizard. Simply accept the default options to complete the installation.

You will likely be prompted to restart the system after the installation completes. *Please allow your system to reboot.*

At this point, with VMware Player installed, the first thing that we will do is install the


virtual machines that we will use. To do this, please look again at the course USB. On the course USB, in the “Days 1-5” folder, you will find the “507 Systems” folder.

Name	Date modified	Type
507 Systems	2/12/2014 5:50 AM	File folder
Checklists	10/13/2013 3:15 PM	File folder
Configs	2/12/2014 5:59 AM	File folder
Extras	12/2/2013 8:41 AM	File folder
Scripts	10/13/2013 3:15 PM	File folder
Standards Docs	10/13/2013 3:15 PM	File folder
Tools	2/11/2014 11:48 AM	File folder
UsefulSpreadsheets	2/11/2014 4:06 AM	File folder
VPN	10/13/2013 3:15 PM	File folder
Information System Accreditation Form...	8/5/2003 9:38 AM	DOC File
VMware-player-6.0.1-1379776	10/20/2013 10:37 ...	Application

After VMWare Player is installed
Please drag the “507 Systems” folder to your system. Where you store the folder is not important, but you should put it in a location where you will be able to easily locate it again when it is required.

This next step is the part of the exercise that will likely irritate your antivirus software! Don’t worry too much about files that cannot be copied. We will deal with these issues as they arise during the week.

Using your Windows file browser, please open the course USB. Locate the “Tools” folder in the “Days 1-5” folder. Please drag the “Tools” folder from the USB to your C drive. You may, of course, put this folder in any other location that you choose; however, the directions in the workbook will expect that the folder is C:\Tools.

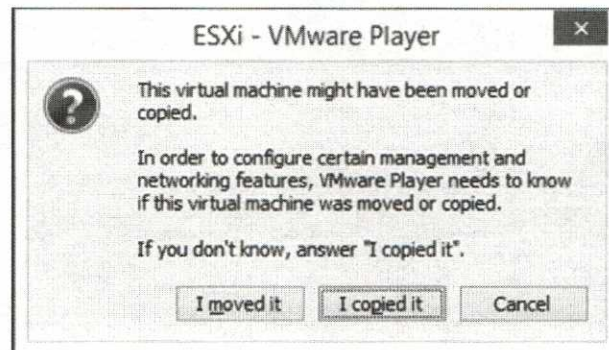


¹ The executable name is “VMware-player-7.1.2-2780323.exe”

Part 1: Starting ESXi

The first step in getting started is to get the ESXi server up and running. You can find the virtual image for this system in the “507 Systems” folder wherever you have placed it on your system. ***It is very important that you do not attempt to run the virtual machines from the USB!*** Although the systems will run, they will run very, very slowly!


After locating the ESXi virtual machine, please double-click the “.vmx” file or configuration file to get it started. When you do, you will be presented with a message something like this:



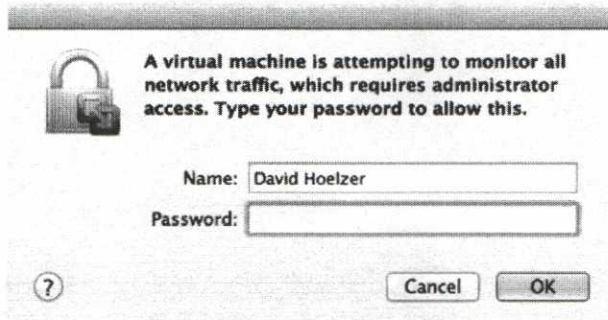
Depending on the version of VMware that you are using, this message might look slightly different, but the question is the same. VMware has noticed that the virtual machine is not in the same location that it was built in and is trying to figure out what’s happened. In our case this week, it does not matter how we answer this question.

The reason for the question is to allow the virtual machine to update its MAC address. ***For today, let’s click “I moved it” because we actually did move it!***

After clicking the “I moved it” button, you should see the virtual machine start up:

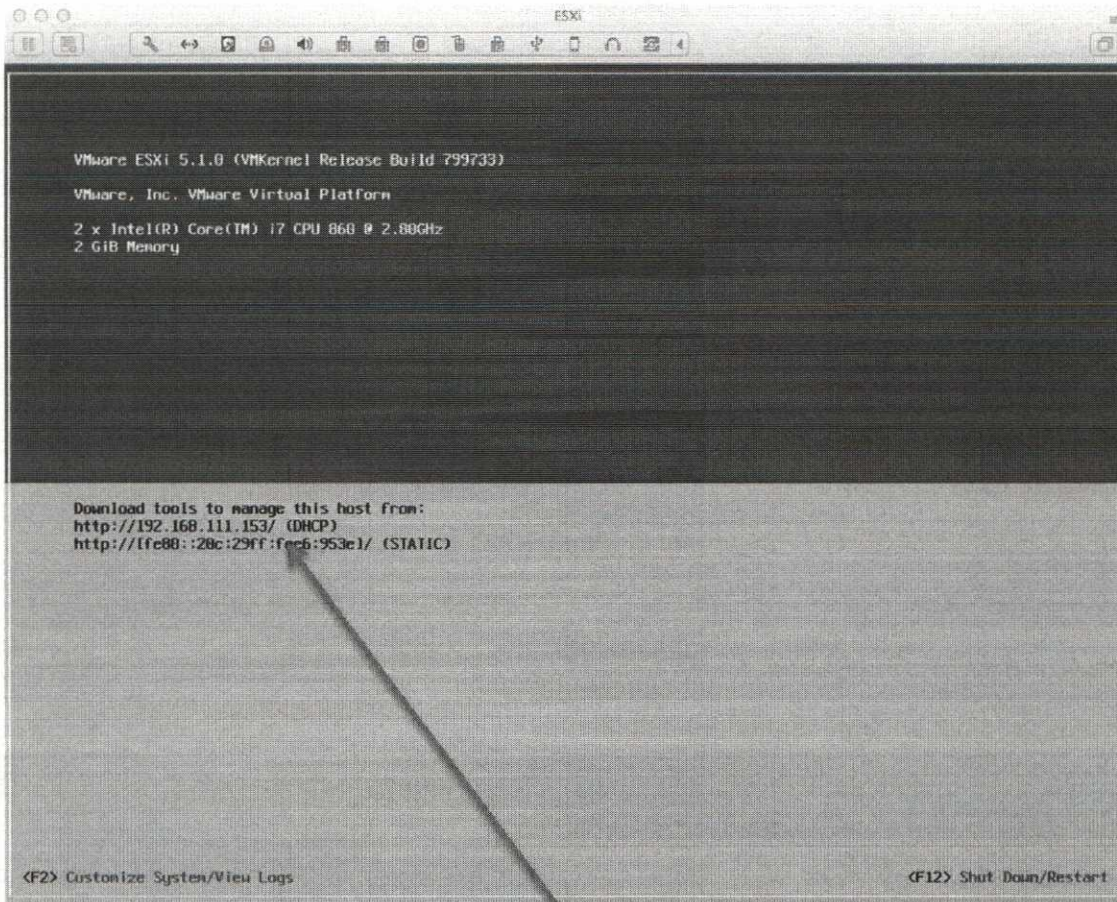


If you are running a Macintosh computer with VMware Fusion, you can also expect the system to prompt you like this:



If your system does prompt you in this way, please enter your password to continue. Because the ESXi system allows you to run multiple virtual instances within it, it needs to run the virtual adapter in promiscuous mode, which your system requires authentication to do.

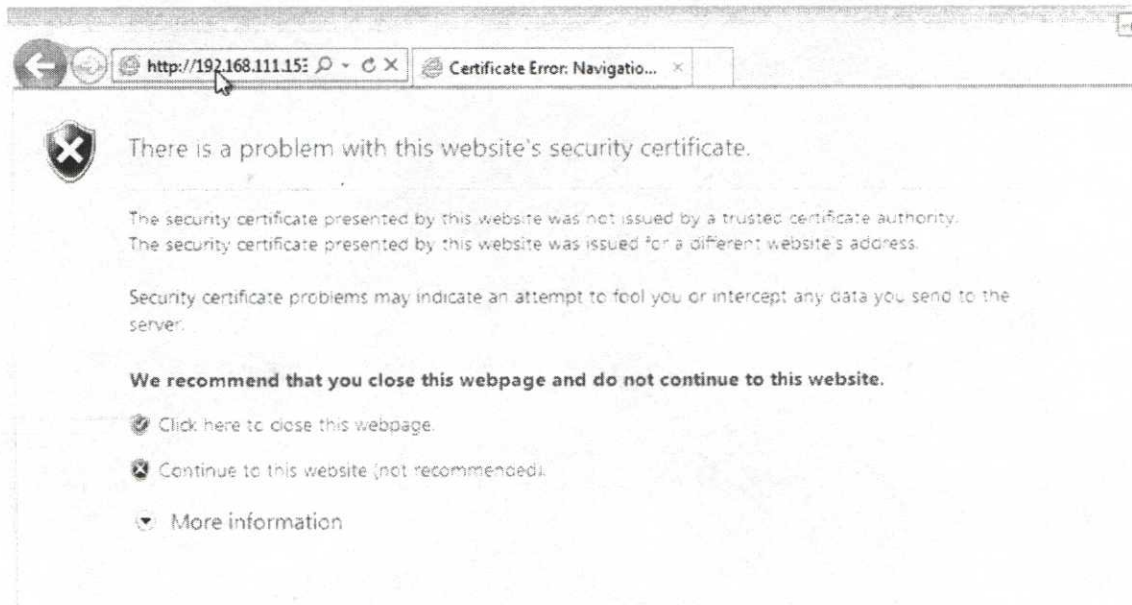
When ESXi finishes booting up, you should see a screen something like this:



Notice that there is a link indicating the current IP address of the virtual system, as indicated. ***Please take note of the IP address that appears in your system. This address is practically guaranteed NOT to match the one in the screenshot above!***

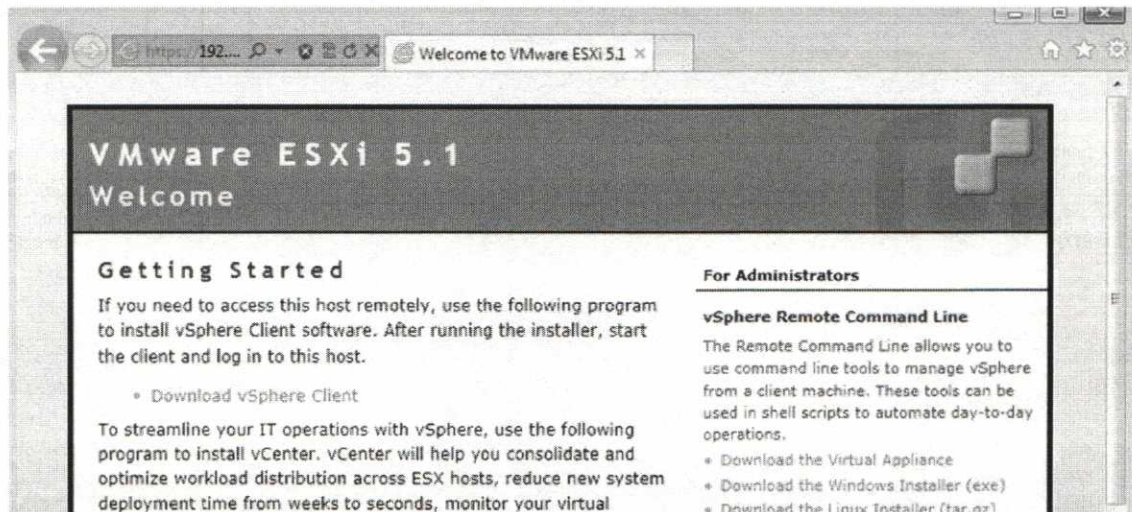
Part 2: Examining the Server

1. Now that everything is up and running, let's explore this system. To do so, ***please start your web browser and browse to the address displayed in the console.*** When you do so, be prepared for an error!



As you can see, your web browser is not comfortable opening this page. What is the reason for this?

2. If you answered, “The certificate isn’t valid,” you’ve answered correctly. It is common to find that the certificates installed on ESXi servers are self-signed certificates. This means that there is no established PKI infrastructure in use, at least within the ESXi infrastructure. ***Why is the use of self-signed certificates a bad practice?***
3. Using self-signed certificates accustoms administrators to bypassing certificate warnings and leaves the organization open to man-in-the-middle attacks using forged certificates. For now, let’s take note of this issue and click through to the website. ***To continue, please note that Internet Explorer’s “continue” option is the one marked with the red shield. Click this button please.*** If you are using the Chrome browser, bypassing this message is not immediately obvious. In Chrome, you must select the “Advanced” option, which gives you the opportunity to bypass or ignore the certificate warning.
4. Clicking through the certificate warning should reveal the following page:



Looking at this screen, is there anything on the page that appears to be particularly sensitive?

5. If you answered, "Browsing datastores in the host's inventory," then you answered correctly! *Please click this link if you haven't already done so.*
6. *Clicking the "Browse datastores in this host's inventory" should cause an authentication prompt to open:*



7. Remember the default username for ESXi systems is "root." In addition, if the ESXi host is anything older than version 5.5, nothing prevents the administrator from choosing a blank password. You might also find that versions 5.5 and above might *still* have a blank password if they were originally ESXi 4 or 5 installs that were later upgraded. *Please enter the default username of "root" with no password and click "OK."*
8. Clicking through gives us instant web-based access to the backing datastore!

Name	Capacity	Free
ESXi Internal Datastore	3221225472	2569011200


What is the potential impact of leaving the datastore unprotected?

What are some effective methods for protecting against web-based datastore browsing?¹

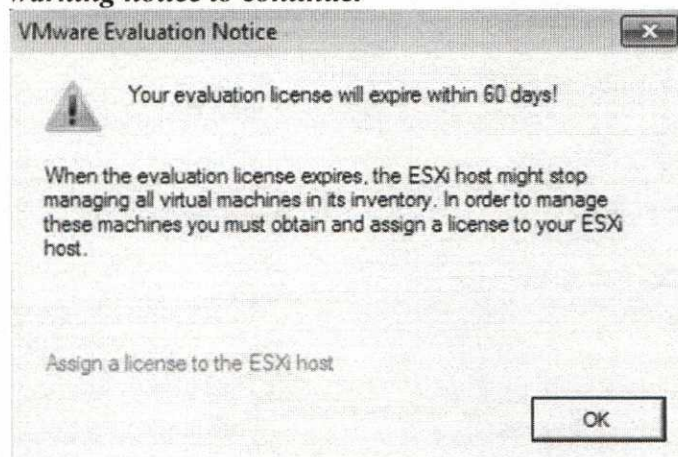
9. At this point, let's get the vSphere client installed so that we can manage the ESXi server directly. Normally, you would click the "Download vSphere Client" link, but this will take you to the VMWare site and can take a significant amount of time to download at a conference. Therefore, ***you will find a copy of the vSphere installer on the USB in the "Tools" folder.*** You can also run the installer from your hard drive provided that you previously copied the "Tools" folder over to the system. ***The full name of the installer that you want to run is "VMware-viclient-all-5.1.0-786111.exe."***
10. ***Please run the installer and accept all of the defaults.***
11. Now that vSphere is installed, let's connect to the ESXi system using this native interface. To do so, ***please start up the vSphere client.***

¹ Possible answers to these questions can be found in the back of the workbook on page 232.

12. When vSphere is first started, you will be presented with a dialog box that looks like this:



13. **Please enter the IP address that is displayed by your ESXi system. Also, please enter the username “root” with no password and click the “Login” button.**
14. You will be prompted regarding the validity of the certificate on the ESXi system. **Please click the “Ignore” button.**
15. After several seconds, you should see the vSphere management interface on your screen with an alert regarding the evaluation status of the server. **Close the warning notice to continue.**



16. Using the vSphere interface, please try to answer the following questions:
- a. **The following questions can be answered using the Summary tab:**
 - i. At least one configuration issue is being brought to your attention. What is(are) this (these) issue(s)?¹
 - ii. What is the current licensing status of this system?
 - iii. How many processor sockets are installed?
 - iv. How many NICs are installed?

¹ Answers to these questions based on the provided ESXi system can be found on page 232.

- v. Are there any potential findings concerning the number of NICs installed?
- vi. How many networks are configured on this ESXi system?
- b. *The following questions can be answered using the Configuration tab:***
 - i.** How much memory is installed on this virtual server?
 - ii.** What is the total capacity of the storage available?
 - iii.** What is the name of the installed datastore?
 - iv.** Is the installed datastore an SSD drive?
 - v.** Which file system is in use on the datastore?
 - vi.** Which networks are connected to the physical network adapter?
 - vii.** What is the VLAN ID of the “Services” network?
 - viii.** Is IPV6 enabled on vmnic0?
 - ix.** To which switch is vmnic0 attached?
 - x.** Is there anything wrong with the time configuration?
 - xi.** Is the system configured to accept domain authentication?
- c. *To answer the following questions, you must click the “Advanced Settings” option under “Software”:***
 - i. Are the system logs configured to be sent to a remote system?
 - ii. What is the IP address of the system?
 - iii. What is the current setting for maximum log size that will cause the logs to be rotated?

Day 2

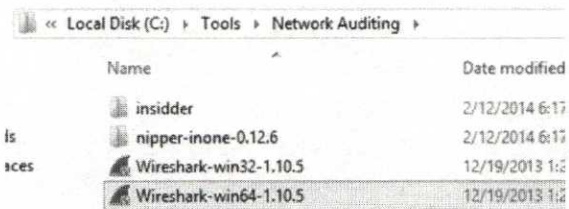
Exercise 1: Switches and Network Symptoms

Time Required: Approximately 30 minutes

Purpose: Familiarize learners with common configuration options and issues. Illustrate easy-to-find evidence indicating incorrect configurations in switched environments. Introduce the Wireshark sniffer.

Part 1: Wireshark – Introduction

To begin, you will learn a little bit about a very handy sniffer named “Wireshark.” It is possible that you have seen this tool in the past before it was renamed. It was previously known as “Ethereal.” Let’s get the sniffer installed so that you can look at some data!



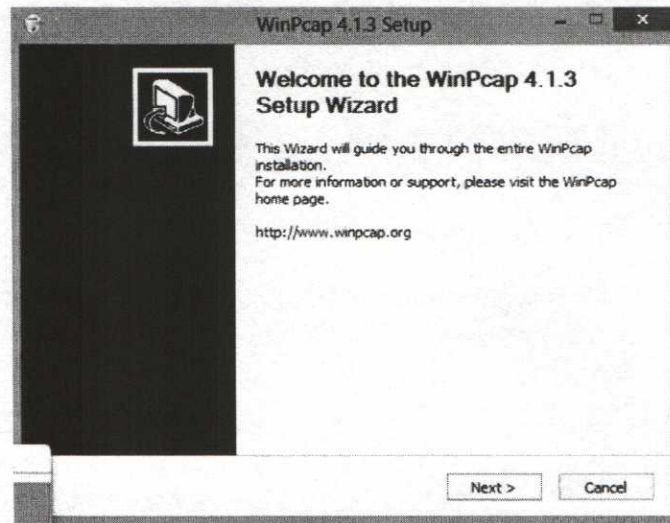
Name	Date modified
insidder	2/12/2014 6:17
nipper-inone-0.12.6	2/12/2014 6:17
Wireshark-win32-1.10.5	12/19/2013 1:2
Wireshark-win64-1.10.5	12/19/2013 1:2

Please locate the Wireshark installer. You will find the installer in the “Tools” folder that you previously copied onto your system. Within that folder, please browse to the “Tools” folder, and then into the “Network Auditing” folder.¹

When running the installer, please accept all defaults.

At some point during the install, a second installer for WinPCAP will open. This is normal and expected! You must also accept all of the defaults for this installer for Wireshark to work properly on your system.

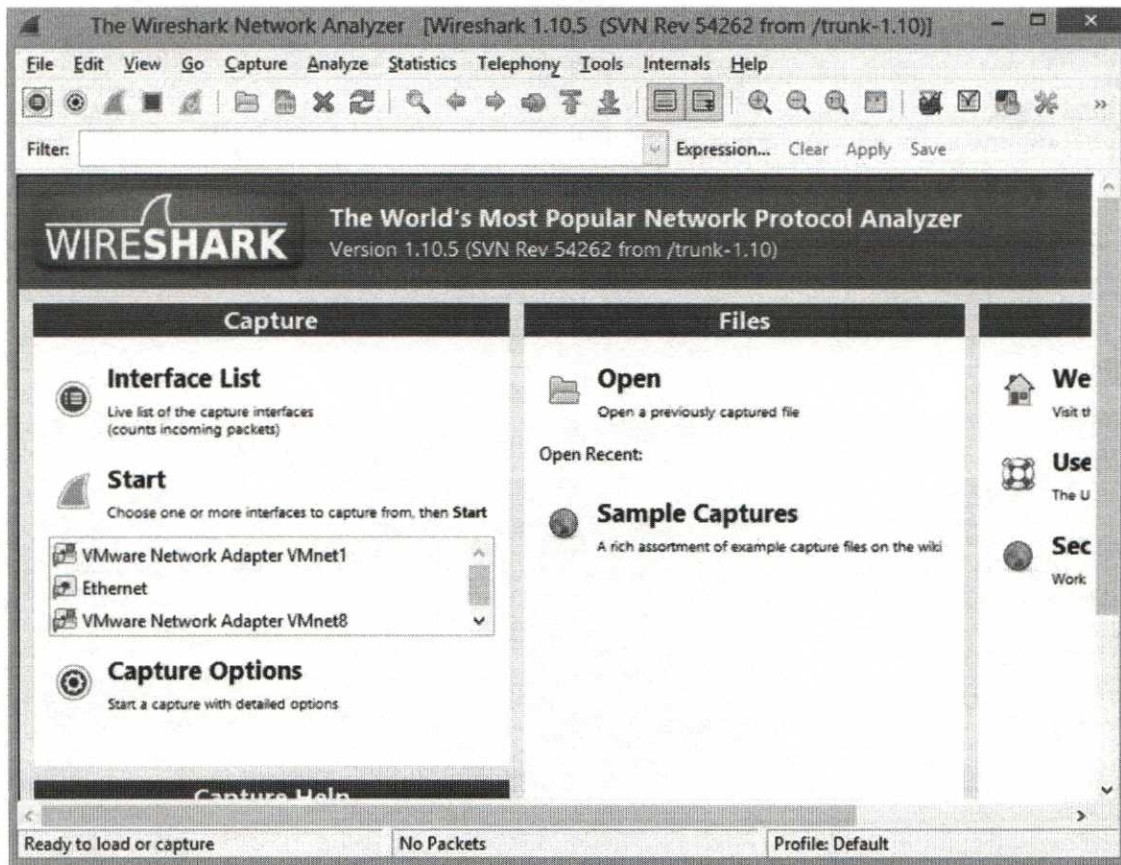
¹ Please note that you might find a slightly newer version of Wireshark on your USB stick. The USB stick tends to be updated more frequently than the workbook.



After the installation completes, you are ready to get started. In this portion of the lab, you look at a packet capture that was made on an internal network in the EnclaveForensics network. In this particular case, the capture was taken from our research lab. Don't worry – our production network is better configured than this!

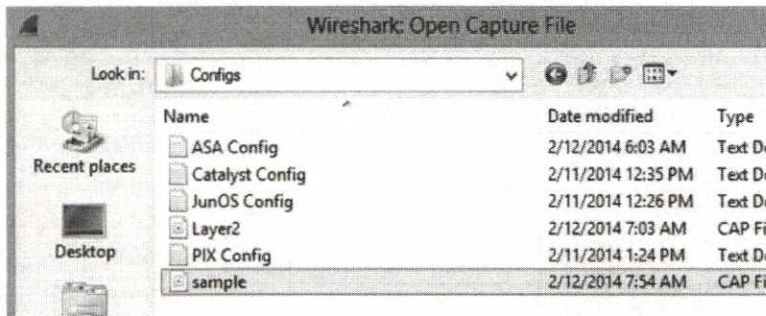
Let's see how easy it is to detect these Layer 2 management protocols that we've mentioned in the class. Let's begin by opening the packet capture and learning a little bit about how Wireshark works.

Please start Wireshark by double-clicking the icon that was placed on your desktop or by locating it via your Start menu or Metro interface.



When Wireshark first starts up, you will be presented with the window that is pictured above. Wireshark, a free open-source product, has become about the best-known sniffer available. It supports hundreds of protocol decodes, allows for network troubleshooting, provides easy mechanisms for reassembling packets into complete conversations, and more.

To get started with this tool, let's first learn about its features and familiarize ourselves with its interface. ***Please click the "Open" option in the main window or pull down the "File" menu and select "Open."*** Using the file browser that appears, please browse to the "Days 1-5" folder on the USB and locate the "Configs" folder. In this folder, find the file named "Sample" and open it.

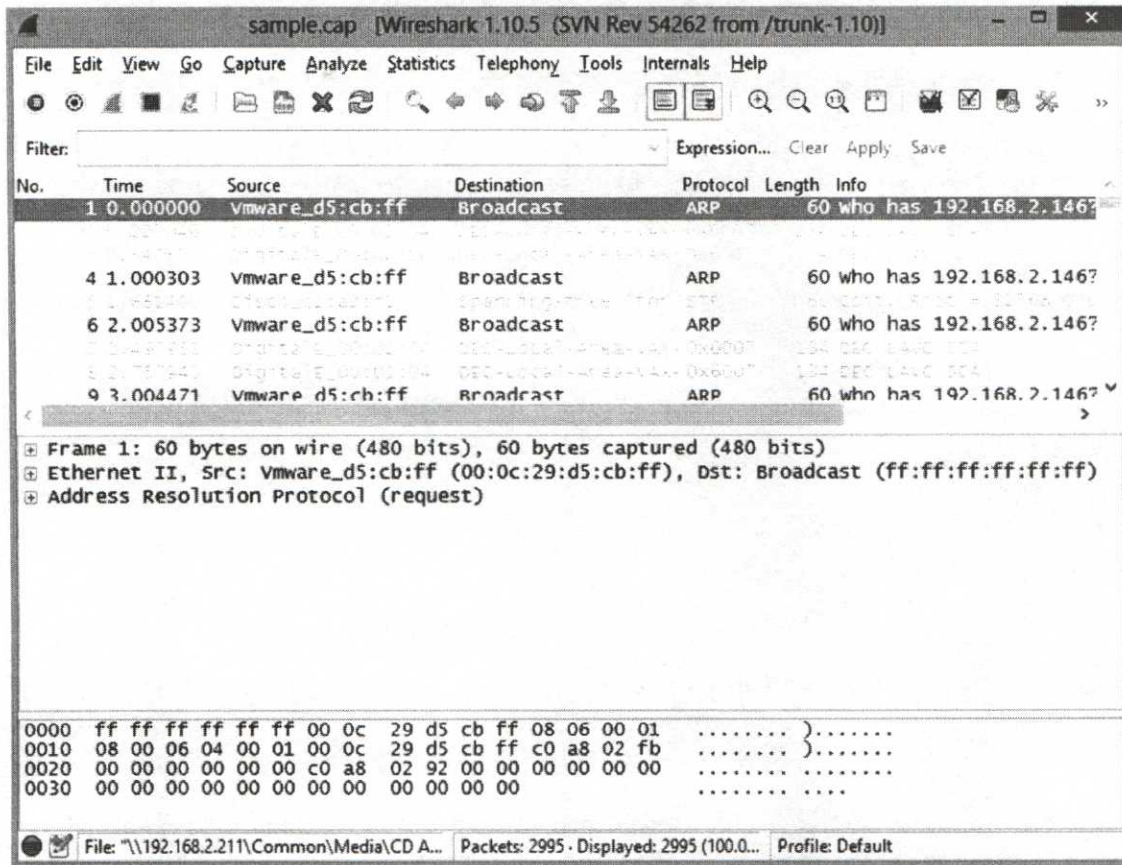


The sample capture that we have included on the USB has a few basic pieces of network data in it that are common to find. We'll use these to familiarize ourselves with the Wireshark interface

and to see how some of the more advanced features work.

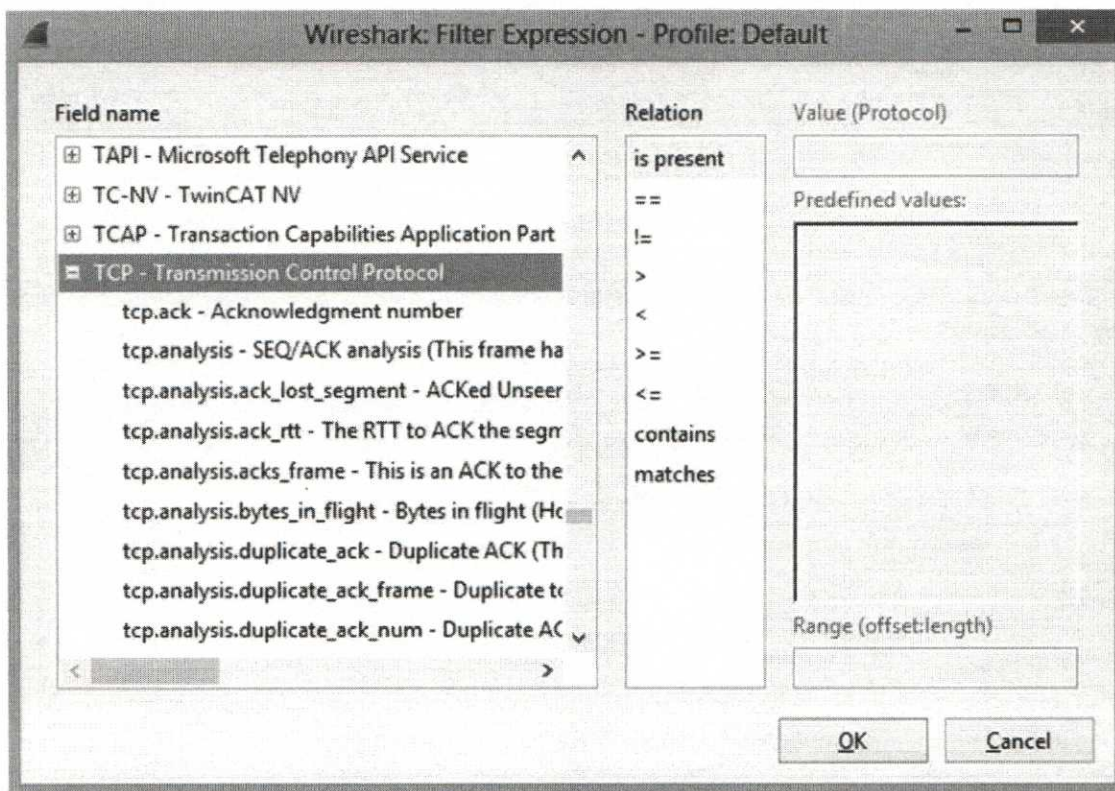
SANS Advanced Systems Audit Workbook

After opening the file, you should see a window that looks like what is pictured below:



The Wireshark interface is fairly easy to understand. At the top, immediately below the menus, are a number of quick links to common features. Immediately below that, you can see a box marked "Filter" with an "Expression" button to the left. This can be used to perform searches or to select which data is displayed in the top window.

Please click the "Expression" button. In the window that opens, please scroll down, locate the "TCP" node, and select it.




The window that you see can be used to interactively build filters or queries to select specific data in the Wireshark main window. Please feel free to experiment here. Note that we look only at TCP traffic for now. ***Now that TCP has been selected, please click the “OK” button to continue. When you return to the main interface, please click the “Apply” button to the right of the expression that you just created.***

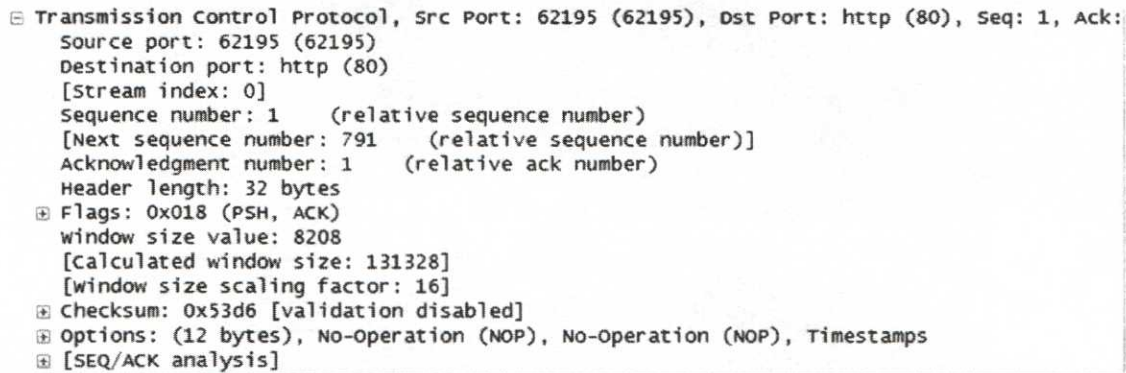
No.	Time	Source	Destination	Protocol	Length	Info
18	4.728243	192.168.2.161	216.34.181.45	TCP	78	62195 > http [SYN] Seq
19	4.765915	216.34.181.45	192.168.2.161	TCP	78	http > 62195 [SYN, ACK
20	4.765996	192.168.2.161	216.34.181.45	TCP	66	62195 > http [ACK] Seq
21	4.766976	192.168.2.161	216.34.181.45	HTTP	856	GET / HTTP/1.1
22	4.803844	216.34.181.45	192.168.2.161	TCP	66	http > 62195 [ACK] Seq
23	4.812450	216.34.181.45	192.168.2.161	TCP	458	[TCP segment of a reas
24	4.812466	216.34.181.45	192.168.2.161	TCP	1434	[TCP segment of a reas
25	4.812577	192.168.2.161	216.34.181.45	TCP	66	62195 > http [ACK] Seq
26	4.812583	192.168.2.161	216.34.181.45	TCP	66	62195 > http [ACK] Seq

When you apply the filter, the top window immediately changes. Notice that the “Protocol” column now reveals that only TCP packets are being displayed. You have likely also noticed that the packets have been color-coded. How data is displayed is completely configurable. We will just use the defaults for now.


This top window that you’re looking at is, in a sense, giving you a sort of “network level” view. It shows you all of the packets, what kind they are, and when they passed by. ***Please use your mouse to click packet #21.***



When you click a packet in the top window, you see that the other windows immediately come to life. The center window allows you to look at a protocol-level decode of the packet. For example, *click the plus sign to the left of the “Transmission Control Protocol” item.*

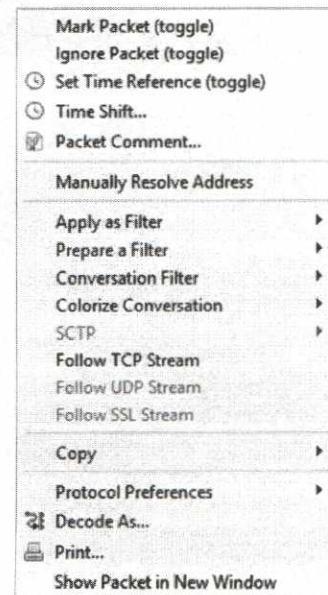


Now that you have opened this section, Wireshark decodes all of the various bits and bytes within the packet that are related to the TCP header, showing you what each piece means. For example, you can see that this is a PUSH-ACK packet coming from port 62195 and going to port 80. Because it is a PUSH-ACK to port 80, you can reasonably assume that this should be some sort of web request. If this were a PUSH-ACK coming from port 80, you might guess that it is a web response. *Please click the “Flags” section in the TCP portion of the display.*



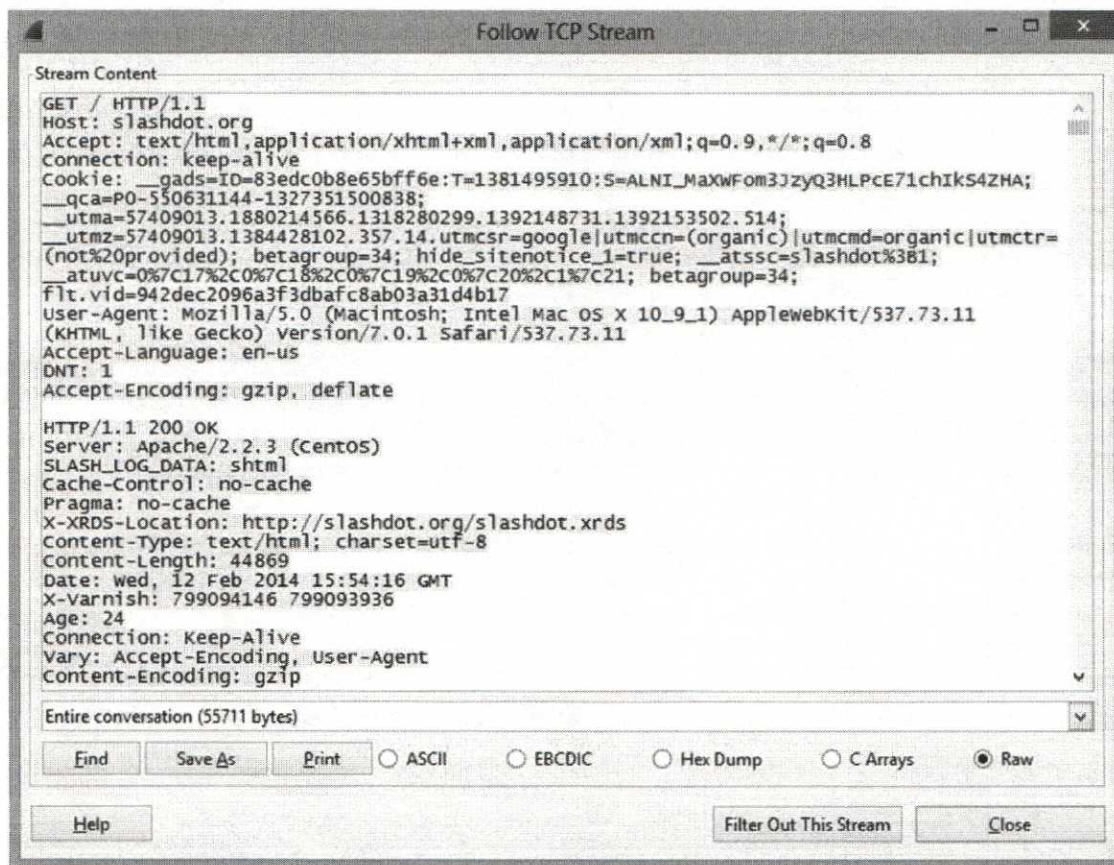
The bottom window in the Wireshark frame displays the actual packet itself. For a network engineer or intrusion analyst, this is a valuable section because it allows the engineer to see for himself exactly what values appear in the packet. Notice that the last two hexadecimal values on the top row are highlighted. *When you select something in the center window, the actual data in the packet will be highlighted for you in the bottom window.*

Let’s look at one other amazing feature that Wireshark has. This is not a feature that you will need right away; however, later



today, we discuss the use of vulnerability scanners and will suggest that you always run a sniffer at the same time. We explain why it is in the course book, but what can we do with the data that is captured? **Please right-click packet #21 in the top window, and select the “Follow TCP Stream” option.**

The “Follow TCP Stream” option analyzes the packet capture and identifies all of the packets that are related to this specific conversation. After the packets are identified, it extracts the content and reassembles all of the payloads to provide a view that shows everything that passed between the client and the server. Because you can reassemble that data, you can also extract that payload. This can be used to recover any arbitrary data from a network stream very, very easily!

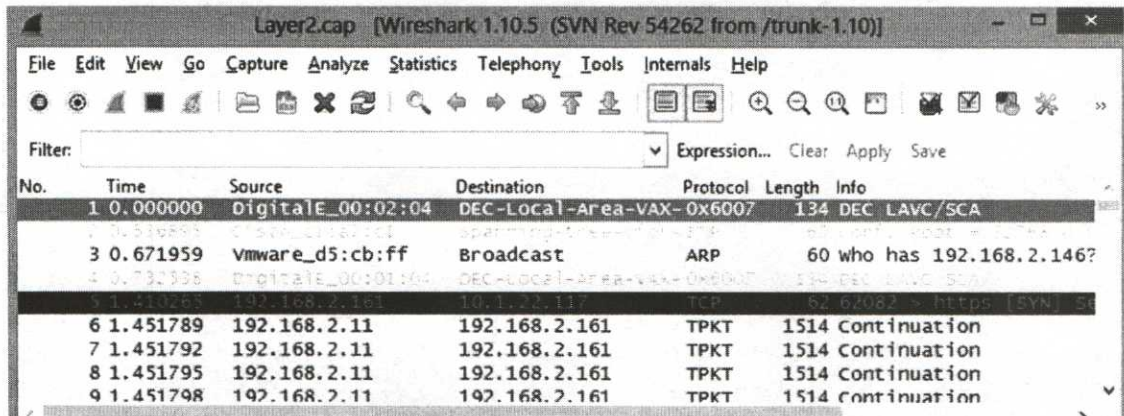


As you can see above, you now have a complete view of everything that has transpired. The data in red represents data that has passed from the client to the server. The data in blue is the data that has returned from the server to the client. Not only can you choose how to display this data, but you can also choose to export this data into a variety of useful formats.

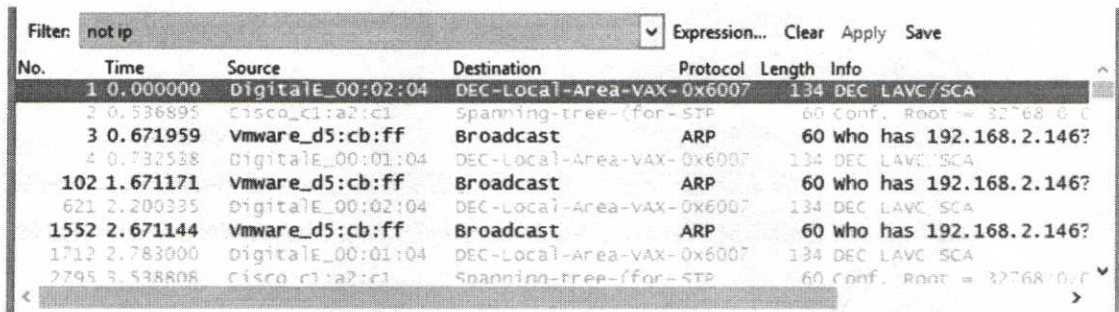
Part 2: Wireshark – Layer 2 Analysis

Now that you are somewhat familiar with Wireshark and some of its features, let's open up a different capture file. **Please close the “sample” capture file. Once again, use the**

“Open” function to locate the capture file named “Layer 2” that is also located in the “Days 1-5” folder on the USB in the “Configs” folder. When you first open this file, it is possible that the previous TCP filter will still be applied. This is the filter that was used to extract the stream from the “sample” capture. **If your view displays only a few TCP packets, please click the “Clear” button to the right of the filter expression.**



In the main window in Wireshark, you can now see the data from the Layer 2 capture. An important exercise that we often conduct with network engineers is to take the time to periodically identify every single type of communication that occurs on the network. Because we’re not training network engineers today, we’re not going to go that far. However, we will use an aspect of this process to look for things that could indicate Layer 2 misconfigurations. **Please type “not ip” into the “Filter” window, and then click the “Apply” button.** As you type, you will notice that the filter window turns red and green. Can you tell what’s happening here? Whenever the value in the filter window is a valid filter, it will turn green. If the expression is not a valid filter, it will turn red to let you know.



After inserting the “not ip” filter and applying it, you can see that all of the TCP and TPKT (another TCP conversation) have been removed from view. Before you analyze, there is another packet type that is visible that you really don’t need to look at. **Notice that ARP packets are still being displayed. ARP is a normal part of IPv4. Although this is certainly operating at Layer 2, it is not the kind of protocol that we’re interested in finding. Please add “and not arp” to the filter expression and click Apply.**

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000	DigitalE_00:02:04	DEC-Local-Area-VAX-0x6007	134	DEC LAVC/SCA	
2	0.536895	Cisco_cl:a2:c1	Spanning-tree-(for-STP	60	Conf. Root = 32768/0/0	
4	0.732538	DigitalE_00:01:04	DEC-Local-Area-VAX-0x6007	134	DEC LAVC/SCA	
621	2.200335	DigitalE_00:02:04	DEC-Local-Area-VAX-0x6007	134	DEC LAVC/SCA	
1712	2.783000	DigitalE_00:01:04	DEC-Local-Area-VAX-0x6007	134	DEC LAVC/SCA	
2795	3.538808	Cisco_cl:a2:c1	Spanning-tree-(for-STP	60	Conf. Root = 32768/0/0	
3265	3.850807	DigitalE_00:02:04	DEC-Local-Area-VAX-0x6007	134	DEC LAVC/SCA	
3414	5.651266	DigitalE_00:02:04	DEC-Local-Area-VAX-0x6007	134	DEC LAVC/SCA	
3492	5.783367	DigitalE_00:01:04	DEC-Local-Area-VAX-0x6007	134	DEC LAVC/SCA	

At this point, you have eliminated all of the obvious “stuff” that can distract you. The only things left should be packets that are not IP and not ARP, which on most networks will leave you with only Layer 2 packets. Let’s see what we actually find.

Select the first packet, packet #1, and break open the Ethernet II header in the center window. What kind of packet does this appear to be? What kind of systems is it for?

Frame 1: 134 bytes on wire (1072 bits), 134 bytes captured (1072 bits)
Ethernet II, Src: DigitalE_00:02:04 (aa:00:04:00:02:04), Dst: DEC-Local-Area-VAX-Cluster-gr
Destination: DEC-Local-Area-VAX-Cluster-groups-SCA_01:01 (ab:00:04:01:01:01)
Source: DigitalE_00:02:04 (aa:00:04:00:02:04)
Type: DEC LAVC/SCA (0x6007)
Data (120 bytes)

The answers to the questions above can be found by looking at the Destination field. Notice that these packets are marked as “DEC-Local-Area-VAX-Cluster” packets. In fact, a VAX cluster generates these packets in the network. They are not IP packets, nor are they ARP, but they are absolutely valid packets for the network. These packets actually serve to demonstrate an interesting side point. Notice that these are “Ethernet II” packets. **Click the packet immediately beneath this, packet #2. What type of Ethernet encapsulation is being used for this packet?**

Frame 2: 60 bytes on wire (480 bits), 60 bytes captured (480 bits)
IEEE 802.3 Ethernet
Logical-Link Control
Spanning Tree Protocol

As you can see, this packet uses 802.3, the most common kind of Ethernet encapsulation in use today on internal networks. Although this is not exactly what we are looking for right now, this is one of the things that we would discuss with network engineers. When you have a network that has existed for a decade or more, it is common to find lots of mysterious packets and encapsulations in use. All of these should be accounted for and, if possible, you should strive to use a single encapsulation type to make the network a bit more efficient.

This second packet, however, is exactly the kind of thing you’re looking for when examining packet captures for Layer 2 issues. **What kind of packet is this?**

The last line in the center window indicates that this is a Spanning Tree packet. As discussed in class, STP, MSTP, TRILL, and SPB (in addition to VTP, 802.1Q, CDP and

other VLAN protocols) are all bad indicators. Switch management protocols and VLAN management protocols indicate that the switches have not been properly locked down. Even though an attacker might not be able to compromise the switch itself, these are strong indicators that he could use a tool such as Yersinia (<http://www.yersinia.net>) or a VLAN stack on a Linux system to subvert any and all Layer 2 controls on the network.

There is another, simpler way to view this type of data at a high level. *Please clear the current filter. After this is done, please pull down the “Statistics” menu and select “Protocol Hierarchy.”*

Wireshark: Protocol Hierarchy Statistics

Display filter: none

Protocol	% Packets	Packets	% Bytes	Bytes	Mbit/s	End Packets	End Bytes
Frame	100.00 %	15313	100.00 %	11235185	2.682	0	
Ethernet	100.00 %	15313	100.00 %	11235185	2.682	0	
Data	0.20 %	31	0.04 %	4080	0.001	31	
Logical-Link Control	0.08 %	13	0.01 %	1097	0.000	0	
Spanning Tree Protocol	0.07 %	11	0.01 %	660	0.000	11	
ISO 9542 ESIS Routeing Information Exchange Protocol	0.01 %	1	0.00 %	60	0.000	1	
Cisco Discovery Protocol	0.01 %	1	0.00 %	377	0.000	1	
Address Resolution Protocol	0.24 %	36	0.02 %	2160	0.001	36	
Internet Protocol Version 4	99.47 %	15232	99.93 %	11227788	2.680	0	
Transmission Control Protocol	99.44 %	15227	99.93 %	11227230	2.680	8088	
User Datagram Protocol	0.03 %	5	0.00 %	558	0.000	0	
DEC DNA Routing Protocol	0.01 %	1	0.00 %	60	0.000	1	

Buttons: Help, Close

As you can see, this gives you a quick-and-easy view of where you can find three kinds of data being broadcast that might be of concern: Spanning Tree Protocol, Cisco Discovery Protocol, and Routing Information Protocol (RIP). We discuss a bit more about RIP and protecting routing paths later today!

In a later lab covering Routers and Firewalls, we will revisit configuration testing with a tool that can identify common configuration errors involving administration interfaces, protocols, etc.

Exercise 2: Network Behavior

Time Required: Approximately 15 minutes


Purpose: Provide hands-on experience with packet crafting tools. Reinforce ability to use a sniffer to answer basic questions. Determine through experimentation how network protocols actually work, allowing for the creation of testing criteria with predictable outcomes.

Preparation: To perform this lab, you need to execute the External server, Internal server, and the Firewall virtual machines. This also requires VMWare Player to be properly installed. If you have not unzipped the virtual machines or installed VMWare Player, please go back and perform the steps outlined on page 11.

Part 1: Virtual Machines

To get started, we need to start up several of the virtual machines that we unpacked yesterday or last night. We will actually use these systems several times today, so make sure that you become familiar with how to start them up!

Please locate the “507 Systems” folder that you copied to your computer yesterday. Within that folder, you should find a “Day 2 VMs” folder.





In the “Day 2 VMs” folder, you will find another folder, “Day 2.” In this folder, there are four folders: DNS, External, Firewall, and Internal.

Open the “DNS” folder and locate the DNS virtual machine configuration file. This file can be easily identified by the three interlocking squares that are used to represent its icon, as illustrated to the right. Double-click this file.



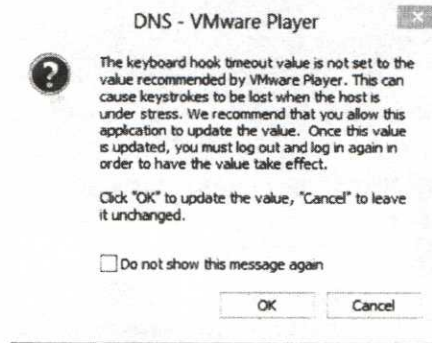
Double-clicking a VMWare configuration file causes VMWare Player to open the virtual machine. If you did not actually meet the laptop requirements and bring a 64-bit operating system, you might receive a warning when opening many of the virtual machines today. For example, you might receive the warning that appears on the right-hand side of the page. This indicates that you are running a 32-bit OS but you are opening a 64-bit OS. There’s nothing you can do about this at this point in the class, so *if you do receive this warning, select the “Do not show this hint again” and click “OK.”*



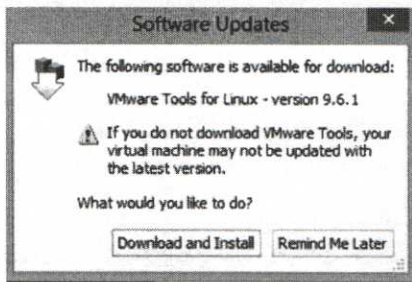
Because the system realizes that it is in a new location, it will ask you to tell it whether, “I moved it” or “I copied it.” For our purposes, the answer to this

question does not matter. Select either option any time you are prompted with this question this week.

After making this prompted with can happen for use this week. We time and expect appropriately in the ***message pictured not show this "OK."***



selection, you might be several other windows. This every virtual machine that we will deal with them just one that you can answer future. ***If you receive the to the right, please select "Do message again" and click***



Although you might not experience the message above, you will almost certainly see the next message at least once this week. VMWare publishes a software package that includes drivers and other modules that allow virtual machines to interact with the host operating system more seamlessly in addition to providing some performance benefits. If we planned to click and drag data into or out of these virtual

machines, these features would be important. For this class, however, you simply use these virtual machines for testing and experimentation. Therefore, you do not need these tools. ***When asked about applying VMWare Tools software updates, please select "Remind Me Later."***

At this point, the DNS server will successfully launch. Whew!

Now that the DNS server is up and running, ***please browse to the "Firewall" folder and follow the same process to start up the Firewall virtual machine.*** The firewall might take slightly longer to start up than the DNS server. Just be patient.

```
Firewall - VMware Player (Non-commercial use only)
Player
*** This is m0n0wall, version 1.8.1
    built on Wed Jan 15 13:32:38 CET 2014 for generic-pc
    Copyright (C) 2002-2014 by Manuel Kasper. All rights reserved.
    Visit http://m0n0.ch/wall for updates.

LAN IP address: 10.17.1.1
WAN IP address: (unknown)

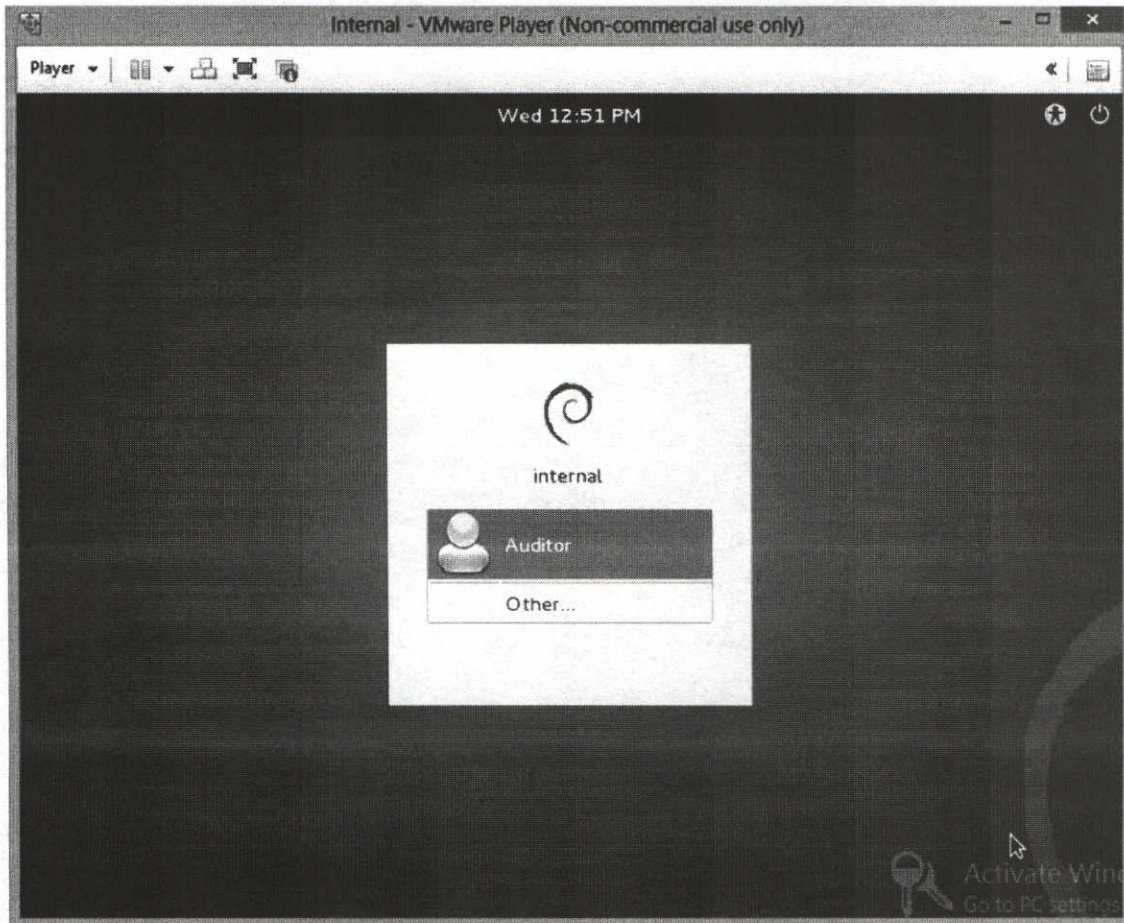
Port configuration:

LAN   -> le1
WAN   -> le0

m0n0wall console setup
*****
1) Interfaces: assign network ports
2) Set up LAN IP address
3) Reset webGUI password
4) Reset to factory defaults
5) Reboot system
6) Ping host

Enter a number: █
```


After the firewall is up and running, *please browse to the Internal folder and start up the Internal virtual system using the same process.*



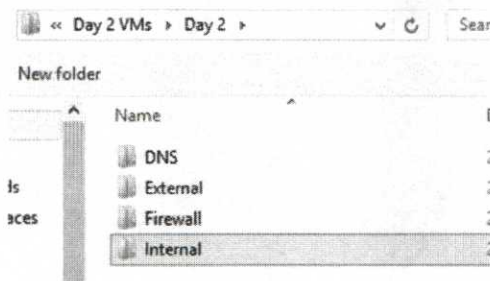
EXTREMELY IMPORTANT!

The External virtual machine must also be started, but there *will* be differences in the startup process!! **Please follow the startup procedures that have been outlined so far with the External virtual machine. We will pick up with the differences below.**

When attempting to start up the External system, you will be informed that an important file could not be located. Please read this message carefully! What is it asking for? Where is this file?

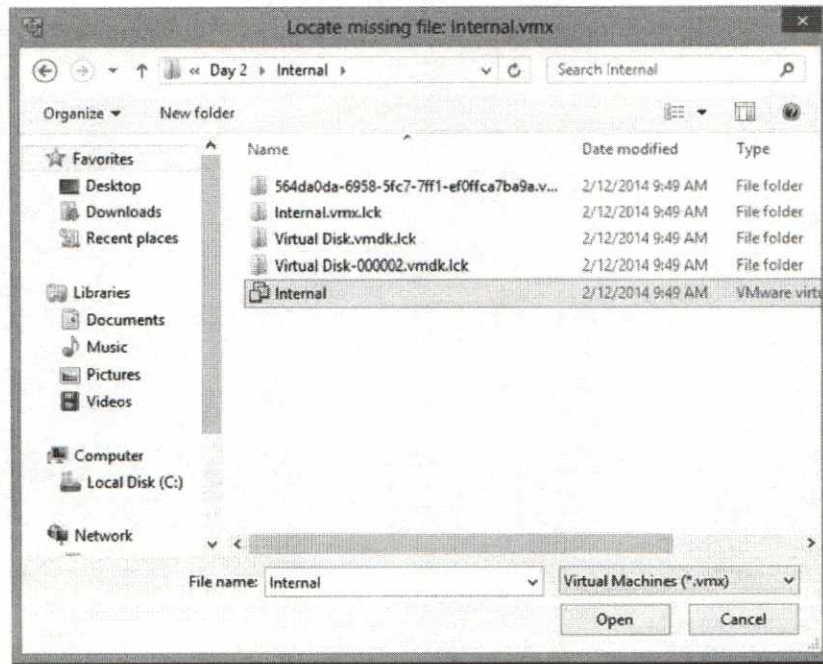


If you read the message carefully, you will see that it is saying that it needs to locate the “Internal.vmwarevm/Internal.vmx” file. Because we have relocated these systems first to a USB and now to your system, the file path has changed. **To resolve this issue, please click the “Browse” link and use the file browser to go to the “Internal” virtual machine folder.**



When you have located the “Internal” folder, double-click it to go into it and find the

“Internal” file with the three blue interlocking squares as an icon. Please select this file and choose “Open.”



EXTREMELY IMPORTANT!

If for some reason your system does not have sufficient RAM, you can accomplish the first lab using the External system only. If you are in this situation, just ask your instructor to clarify precisely which systems are necessary during each lab.

At this point, you should have successfully started up all of the virtual systems that we will use today. **You will not need to log into either the firewall or the DNS server!** In fact, we will not be providing you with credentials to these two systems.

On the other hand, you will use the Internal and External systems extensively.

LOGON CREDENTIALS:

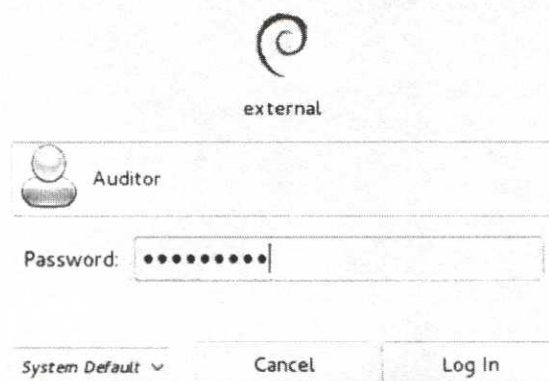
Internal and External have the same credentials.

Username: **auditor**

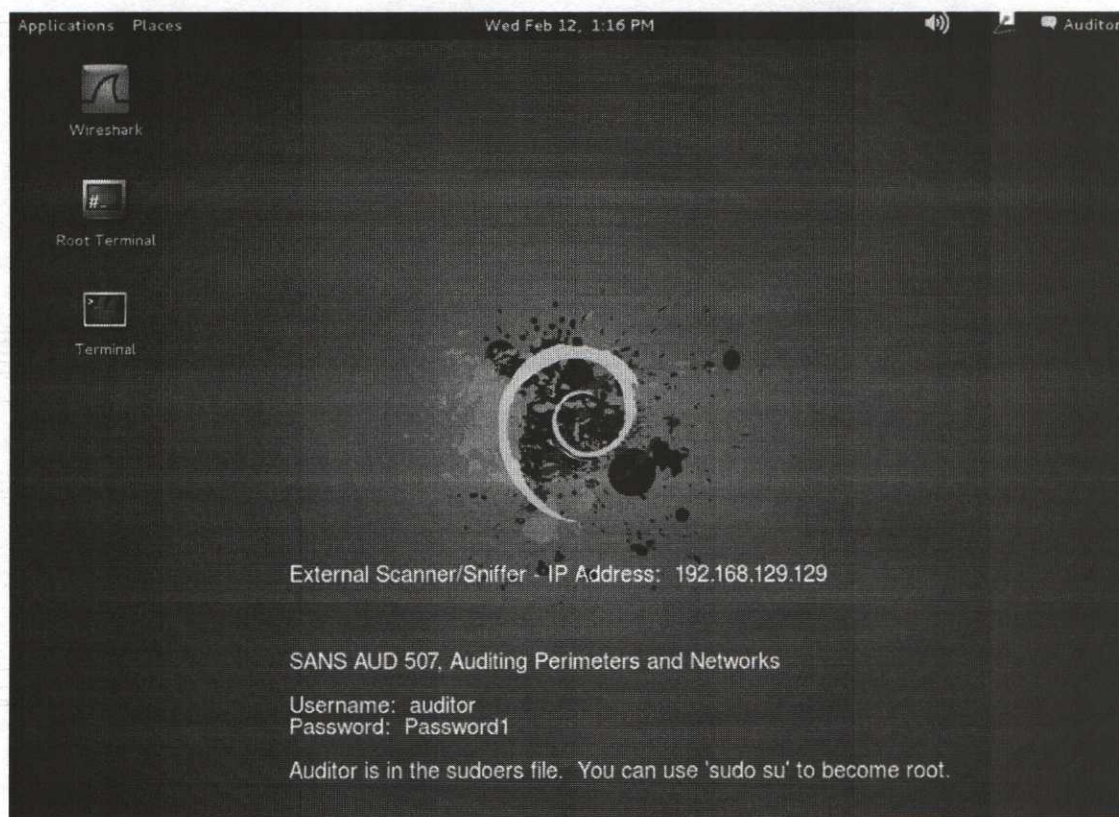
Password: **Password1**

Please click the “External” system and click the Auditor name. Please log in.

SANS Advanced Systems Audit Workbook



After logging in, you will see the user desktop, pictured below.



On the left side of the desktop, there are several icons for tools that we will use today. These include Wireshark, a link to a “root” terminal, and a normal terminal that will run as “auditor.”

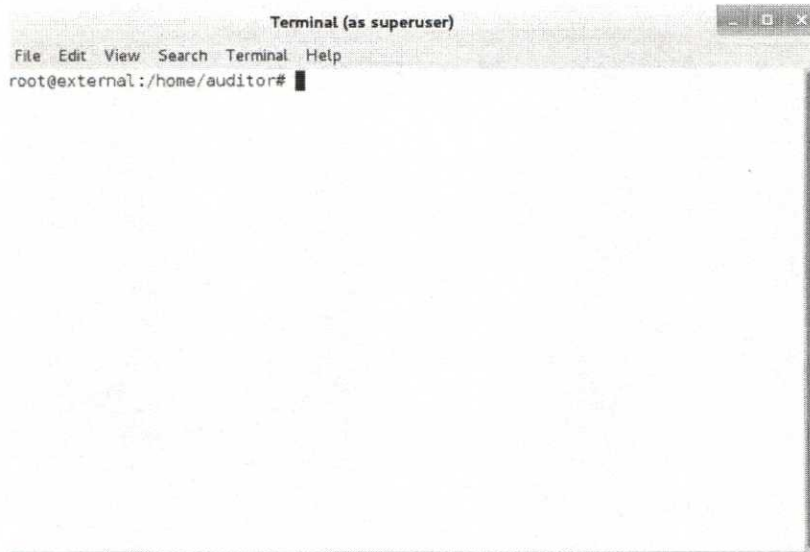
Also, notice that information about the External virtual machine is displayed on the desktop. **PLEASE NOTE! The IP address that your External machine reports will almost certainly be different from what is pictured in the workbook. Please take note of what your IP address is reported as.**

SANS Advanced Systems Audit Workbook

The reason that this happens is that the External system and the Firewall virtual machine are configured to run as “NAT” systems. This means that if you want to use the External machine later this week to scan something else, for example, you would be able to do so without doing any reconfiguration.

On the other hand, the DNS and Internal systems and the internal interface of the firewall are statically configured. We provide you with a network diagram in a later exercise.

For now, let’s get started with our experiment. ***Please double-click the “Root Terminal” icon.***



The root terminal icon provides you with a quick way to start up a terminal window running as the root user. This is required because you are about to use HPing to spoof packets. Spoofing packets requires root privileges in a UNIX environment because you will access the network adapter directly, creating packets that meet your

specifications. We’d like to use the HPing tool to see how a system responds to several different tests.

Please run “hping3 -h | more” in the root command shell window. Use the space bar to page through the output. Please answer the following questions:

Which option allows you to send a TCP packet?

Which option allows you to send a SYN packet?

Which option allows you to turn on the ACK bit in a packet?

Which option allows you to turn on the FIN bit in a packet?

Which option specifies the destination port number for the packet?

For your convenience, the output of the “-h” option is included in the following:

```
usage: hping3 host [options]
```

SANS Advanced Systems Audit Workbook

```
-h --help          show this help
-v --version      show version
-c --count        packet count
-i --interval     wait (uX for X microseconds, for example -i u1000)
                  --fast      alias for -i u10000 (10 packets for second)
                  --faster    alias for -i u1000 (100 packets for second)
                  --flood     sent packets as fast as possible. Don't show
replies.
-n --numeric      numeric output
-q --quiet        quiet
-I --interface    interface name (otherwise default routing interface)
-V --verbose      verbose mode
-D --debug        debugging info
-z --bind         bind ctrl+z to ttl(default to dst port)
-Z --unbind       unbind ctrl+z
--beep           beep for every matching packet received

Mode
default mode     TCP
-0 --rawip       RAW IP mode
-1 --icmp        ICMP mode
-2 --udp         UDP mode
-8 --scan        SCAN mode.
-9 --listen      listen mode

IP
-a --spooft      spoof source address
--rand-dest      random destination address mode. see the man.
--rand-source    random source address mode. see the man.
-t --ttl         ttl (default 64)
-N --id          id (default random)
-W --winid       use win* id byte ordering
-r --rel         relativize id field (to estimate host traffic)
-f --frag        split packets in more frag.(may pass weak acl)
-x --morefrag    set more fragments flag
-y --dontfrag    set dont fragment flag
-g --fragoff     set the fragment offset
-m --mtu         set virtual mtu, implies --frag if packet size > mtu
-o --tos         type of service (default 0x00), try --tos help
-G --rroute      includes RECORD_ROUTE option and display the
route buffer
--lsrr          loose source routing and record route
--ssrr          strict source routing and record route
-H --ipproto     set the IP protocol field, only in RAW
IP mode

ICMP
-C --icmptype    icmp type (default echo request)
-K --icmpcode    icmp code (default 0)
--force-icmp     send all icmp types (default send only
supported types)
--icmp-gw        set gateway address for ICMP redirect
(default 0.0.0.0)
--icmp-ts        Alias for --icmp --icmptype 13 (ICMP timestamp)
--icmp-addr      Alias for --icmp --icmptype 17 (ICMP address
subnet mask)
--icmp-help      display help for others icmp options

UDP/TCP
-s --baseport    base source port (default random)
-p --destport    [+] [+]<port> destination port(default 0)
```

SANS Advanced Systems Audit Workbook

```
-k --keep          ctrl+z inc/dec
                    keep still source port
-w --win          wintsize (default 64)
-O --tcpoff       set fake tcp data offset (instead of tcphdrln / 4)
-Q --seqnum       shows only tcp sequence number
-b --badcksum     try to) send packets with a bad IP checksum many
                    systems will fix the IP checksum sending the packet
                    so you'll get bad UDP/TCP checksum instead.

-M --setseq       set TCP sequence number
-L --setack       set TCP ack
-F --fin          set FIN flag
-S --syn          set SYN flag
-R --rst          set RST flag
-P --push         set PUSH flag
-A --ack          set ACK flag
-U --urg          set URG flag
-X --xmas         set X unused flag (0x40)
-Y --ymas         set Y unused flag (0x80)
--tcpexitcode     use last tcp->th_flags as exit code
--tcp-timestamp   enable the TCP timestamp option to guess the
HZ/uptime
Common
-d --data         data size (default is 0)
-E --file         data from file
-e --sign         add 'signature'
-j --dump         dump packets in hex
-J --print        dump printable characters
-B --safe         enable 'safe' protocol
-u --end          tell you when --file reached EOF and prevent rewind
-T --traceroute   traceroute mode (implies --bind and --ttl 1)
--tr-stop         Exit when receive the first not
                    ICMP in traceroute mode
--tr-keep-ttl     Keep the source TTL fixed, useful to monitor
                    just one hop
--tr-no-rtt       Don't calculate/show RTT information in
                    traceroute mode
ARS packet description (new, unstable)
--apd-send        Send the packet described with APD
                    (see docs/APD.txt)
```

With that little bit of research done, let's try to actually send some packets and see how the host responds. Let's start by just running HPing with a destination host specified. ***Please run HPing, sending a packet to your own External host with no other options specified. Allow this to run for a few seconds, and then press CTRL-C. Please examine the output. What is it showing you? What happened?***

```
root@external:/home/auditor# hping3 192.168.129.129
HPING 192.168.129.129 (eth0 192.168.129.129): NO FLAGS are set, 40 headers + 0 data bytes
^C
--- 192.168.129.129 hping statistic ---
5 packets transmitted, 0 packets received, 100% packet loss
round-trip min/avg/max = 0.0/0.0/0.0 ms
```

As you can see, you received no responses at all. According to the information that HPing reports, it send packets with "NO FLAGS" set. Here, it refers to the TCP flags that

were set in the packets that were sent. It also tells you that the packets were 40 bytes in size, headers only, with no data included.

Can you guess why we received no responses?

If you think about it, this makes sense. Because there were no TCP code bits enabled, there isn't an appropriate answer.

Please create a command line that will send a SYN packet to port #80. Port 80 on the External host is closed. What is the command line and how does the host react?

```
root@external:/home/auditor# hping3 -S -p 80 192.168.129.129
HPING 192.168.129.129 (eth0 192.168.129.129): S set, 40 headers + 0 data bytes
^C
--- 192.168.129.129 hping statistic ---
4 packets transmitted, 0 packets received, 100% packet loss
round-trip min/avg/max = 0.0/0.0/0.0 ms
```

In this case, lack of response is not as easy to explain away! If you are familiar with TCP, then you know that there should be *something* sent back! ***Can you think of a reason that there is still no response from your External host?***

A little pondering will probably lead you to conclude that there could be a host-based firewall interfering with what you see. Another possibility is that there is some peculiarity being introduced by the use of virtual networking interfaces. Can you get around that? Absolutely! The easiest way to do so is to simply try to send your tests to the localhost address or to a physical system that you know is reachable. Because the localhost address also eliminates the possibility of firewall interference, use that option.

Please repeat the first two tests again, this time targeting the localhost address. How does the host react now? What response do you receive to a SYN packet to a closed port?

You can see HPing with no options sent to localhost:

```
root@external:/home/auditor# hping3 localhost
HPING localhost (lo 127.0.0.1): NO FLAGS are set, 40 headers + 0 data bytes
len=40 ip=127.0.0.1 ttl=64 DF id=18182 sport=0 flags=RA seq=0 win=0 rtt=0.2 ms
len=40 ip=127.0.0.1 ttl=64 DF id=18183 sport=0 flags=RA seq=1 win=0 rtt=0.2 ms
len=40 ip=127.0.0.1 ttl=64 DF id=18184 sport=0 flags=RA seq=2 win=0 rtt=0.2 ms
len=40 ip=127.0.0.1 ttl=64 DF id=18185 sport=0 flags=RA seq=3 win=0 rtt=0.3 ms
^C
--- localhost hping statistic ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max = 0.2/0.3/0.3 ms
```

SANS Advanced Systems Audit Workbook

Below we can see HPing run again, this time targeting port 80 with a SYN:

```
root@external:/home/auditor# hping3 -S -p 80 localhost
HPING localhost (lo 127.0.0.1): S set, 40 headers + 0 data bytes
len=40 ip=127.0.0.1 ttl=64 DF id=18186 sport=80 flags=RA seq=0 win=0 rtt=0.2 ms
len=40 ip=127.0.0.1 ttl=64 DF id=18187 sport=80 flags=RA seq=1 win=0 rtt=0.6 ms
len=40 ip=127.0.0.1 ttl=64 DF id=18188 sport=80 flags=RA seq=2 win=0 rtt=0.4 ms
len=40 ip=127.0.0.1 ttl=64 DF id=18189 sport=80 flags=RA seq=3 win=0 rtt=0.2 ms
^C
--- localhost hping statistic ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max = 0.2/0.3/0.6 ms
```

Based on what you have seen so far, you can see that a SYN to a closed port responds with a RST-ACK. Of interest, please notice that HPing with no options shows a packet returning *from* port zero! In class, you will hear us say that port zero is legal though not used. Here's your proof!

Please write command lines to send an ACK to an open port, a SYN/FIN to an open port, and no options to an open port. You can use port 111 as an open port on this system.

How does the host respond in each case?

ACK to an open port:

SYN/FIN to an open port:

No options to an open port:

Do any of these answers surprise you?

ACK to an open port:

```
root@external:/home/auditor# hping3 -A -p 111 localhost
HPING localhost (lo 127.0.0.1): A set, 40 headers + 0 data bytes
len=40 ip=127.0.0.1 ttl=64 DF id=18197 sport=111 flags=R seq=0 win=0 rtt=0.7 ms
len=40 ip=127.0.0.1 ttl=64 DF id=18198 sport=111 flags=R seq=1 win=0 rtt=0.6 ms
```

SYN/FIN to an open port:

```
root@external:/home/auditor# hping3 -S -F -p 111 localhost
HPING localhost (lo 127.0.0.1): SF set, 40 headers + 0 data bytes
^C
--- localhost hping statistic ---
4 packets transmitted, 0 packets received, 100% packet loss
```

No options to an open port:

```
root@external:/home/auditor# hping3 -p 111 localhost
HPING localhost (lo 127.0.0.1): NO FLAGS are set, 40 headers + 0 data bytes
^C
--- localhost hping statistic ---
2 packets transmitted, 0 packets received, 100% packet loss
```

An interesting thing to note is that the system does not respond to a SYN/FIN. You will find that some hosts will respond with a SYN/ACK!!

Exercise 3: Routers, Firewalls, and Validation

Time Required: Approximately 45 minutes

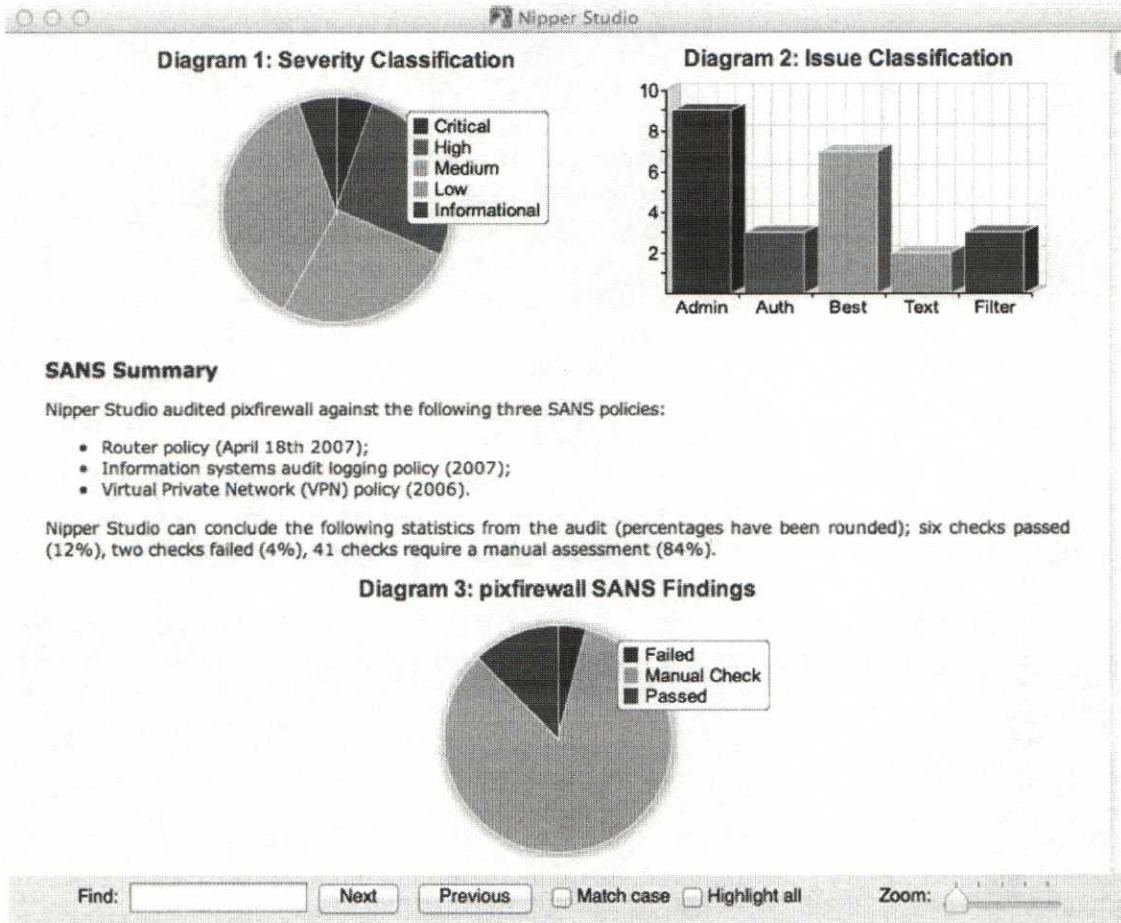
Purpose: Introduce Nipper tool as an example of configuration validation and analysis. Compare reports from switches, routers, and firewalls. Provide hands-on experience performing firewall validations manually. Illustrate use of automated firewall validation scripts.

Preparation: The second part of this lab requires the External, Internal, and Firewall virtual machines to be running. If you have not yet followed the directions for installing and running these systems, please go back and do so now.

Part 1: Nipper

Nipper is a commercial tool that can be used for analyzing firewall, router, and switch configurations. The tool itself supports a wide range of systems including Cisco Catalyst, IOS, ASA and PIX, Netscreen and JunOS devices, 3COM, Bay/Nortel, Checkpoint, F5, and more. Pretty much any network device that makes use of a text-based configuration can be analyzed by this extremely useful system.

The commercial tool is licensed based on the number of nodes and overall size of your routing/switching infrastructure. If you are just trying to look at your perimeter, you can probably get away for about \$1,000 per year. You might feel that that price seems high, but wait until you see what this tool can do, and you're just using the old free version!



Above you can see the high-level graphs that are produced with the commercial version of this tool. Let's dig in with the free version and see what you can find. Along the way, you'll agree that a tool such as this is invaluable in identifying general configuration issues.

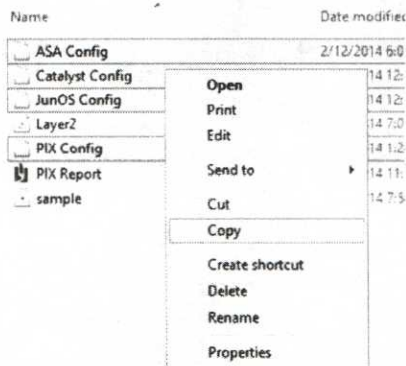
These tools also provide support for validation of ACLs on routers and firewalls. Although you can certainly use these to examine the rules as well, we never do so. In security practice, it is easier in the long term to perform a thorough audit the first time and then verify that any changes to *that* configuration have corresponding entries in the change control system. In addition, because you will validate the firewall technically, you will know whether or not the rules are working properly!

Please locate the "Network Auditing" folder in the "Tools" folder.

Earlier in the workbook, you were asked to copy the "Tools" folder onto your computer. Specifically, we asked that you copy into the root of the C drive so that it would be easy to locate.

In a second Windows Explorer window, please find the "Configs" folder within the "Days 1-5" folder on the USB. Please open this folder.

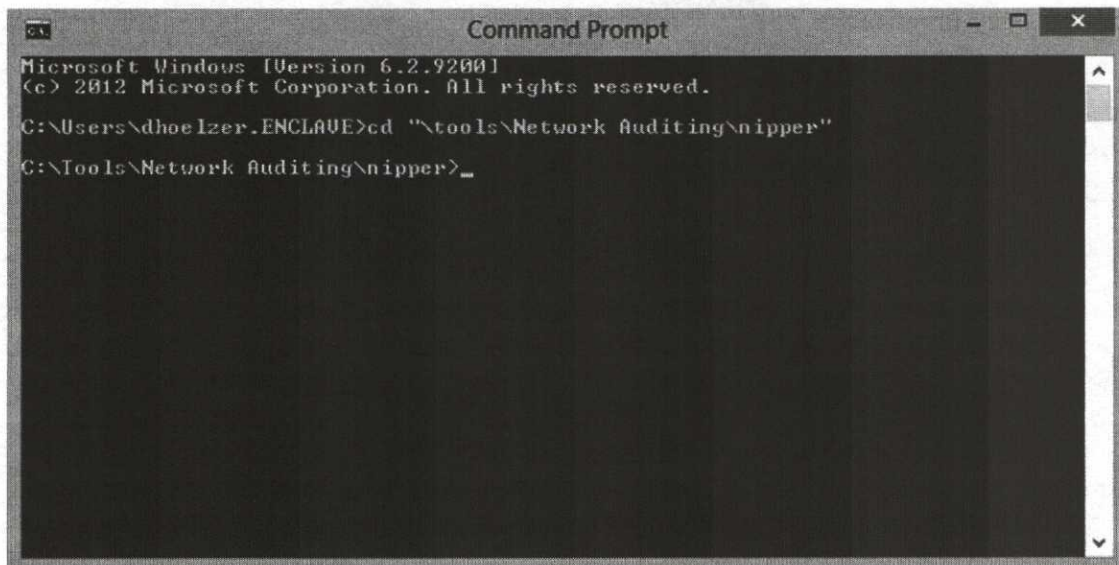
The two sample packet captures from earlier today are in the “Configs” folder. You will also find a sample report from Nipper Studio (PIX Report) and four configuration files.



Please copy the four configuration files, “ASA Config,” “Catalyst Config,” “JunOS Config,” and “PIX Config” to the Nipper folder.

One downside of the free version of Nipper is that it is exclusively a command-line tool. This isn’t a big downside, but this is why we need to get the configuration files into a convenient location to do the analysis of them.

After copying these files, please open a command prompt and navigate to the “C:\Tools\Network Auditing\Nipper” directory.



From this location, you can run Nipper easily and have access to the configuration files. Let’s start by running Nipper against the configuration from a Catalyst switch. This will allow you to tie up any loose ends resulting from the Switch material.

Run the Nipper tool with the “--help” option. Which options are required to specify the source configuration file and the destination where the report will be stored?

As you can see, there are only a few high-level options available, but it is possible to get more details. Notice, particularly, that the help option takes additional options, allowing you to find out what kinds of devices are supported, specify ACL configurations and more.

For now, the options that you're looking for are the "--input=" and the "--output=" options. *Please run nipper using the help feature to determine which kinds of devices are supported.*

Running Nipper with the "--help=devices" option reveals a list of common kinds of systems that we might find on our networks:

```
C:\Tools\Network Auditing\nipper>nipper --help=devices
          O
      O   O   O   O   O   O   O
     O O O O O O O O O O O O
    O O O O O O O O O O O O
   O O O O O O O O O O O O
  O O O O O O O O O O O O
 O O O O O O O O O O O O
O O O O O O O O O O O O

CLI Version 0.12.0
http://nipper.titania.co.uk
Copyright (C) 2006-2008 Ian Ventura-Whiting

Nipper supports a number of different types of network device. This
version contains support for the following devices:

  CMD Option      Device Type
  -----
--auto            Auto-Detect Device (Default)
--3com-firewall   3Com SuperStack 3 Firewall
--accelar         Bay Networks Accelar
--cp-firewall     CheckPoint Firewall Module
--cp-management  CheckPoint Management Module
--ios-router      Cisco IOS-based Router
--ios-catalyst   Cisco IOS-based Catalyst Switch
--pix            Cisco PIX-based Firewall
--asa            Cisco ASA-based Firewall
--fwsn           Cisco FWSM-based Router
--catos         Cisco CatOS-based Catalyst
--nmp           Cisco NMP-based Catalyst
--css           Cisco Content Services Switch
--procurve      HP ProCurve Switches
--screenos      Juniper NetScreen Firewall
--nokiaip       Nokia IP Firewall
--passport      Nortel Passport Device
--nortel-switch Nortel Ethernet Routing Switch 8300
--sonicos       SonicWall SonicOS Firewall

For additional help:
--help[=<topic>]
Show the online help or show the additional help on the topic
specified. The help topics are: GENERAL, DEVICES, DEVICES-ADU,
SNMP, REPORT, REPORT-ADU, REPORT-SECT, REPORT-HTML, REPORT-LATEX,
AUDIT-ACL, AUDIT-PASS, AUDIT-ADU or CONFIG-FILE.
```

Please run Nipper using the PIX configuration file, creating a report named "PIX.html" (Hint: nipper --input="PIX Config.txt" --output="PIX.html" -pix)

When you run Nipper, you might be surprised at how quickly it runs. In fact, most people think that it must not have worked properly because it returns so quickly with no message! *At your command prompt, please type "PIX.html" to open the report in your default web browser.*

Please browse through the report and see whether you can answer these questions:

There is a user named "Admin" configured. What is this user's password? Which privilege level does this password have access to?

On Cisco devices, there are 15 privilege levels used. When you first log on, you are at the lowest. When you type the "enable" command to enable configuration and other administrative functions, you are promoted to level 15.

How many ACLs have been configured on this PIX device?

The answer to this question might surprise you. We would like to point out that all four of these configurations are sanitized examples of actual devices that we have looked at on customer networks.

Which remote administration systems are enabled? Are there any issues with this?

Section 2.4 points out that Telnet is enabled. This is clearly bad, especially for a firewall, but also for any switch or router. We also find that SSH version 1 is enabled. This is bad for a number of reasons. Not only is this version of the protocol vulnerable to man-in-the-middle attacks, but it severely limits the overall key length that can be used to secure the communication. In section 2.14, you also find that SNMP is enabled with a default community string.

What is the logon banner that will be displayed?

Section 2.10 points out that there is actually no logon banner configured at all! Every organization has some kind of requirement with regard to logon banners. You must verify that the logon banner complies with those requirements.

Another important and related setting is in section 2.15. Here, you find that there is also no "Post" logon banner. This is more often known as the Message of the Day (MOTD). It is advisable to synchronize these two messages because some SSH clients will fail to display the logon banner ahead of authentication. In this way, you can be sure that someone authenticating to the device has in fact seen the warning banner.

Please use Nipper to analyze the configuration from the Catalyst switch. Use the generated report to answer the following questions:

What is the configured hostname for this switch?

What is the password used to log onto this system remotely?

There is a Read/Write community string configured. What is it?

Is Telnet supported?

Has remote administrative access been limited to only hosts that should be administering this switch?

Which ports are configured to support trunking? Is this bad? If yes, why?

Are there any other Layer 2 management type services enabled?

Is the HTTP administration interface enabled? If it is, why is this bad?

What is the logon banner for this switch configured to say?

Are the administrative passwords stored securely?

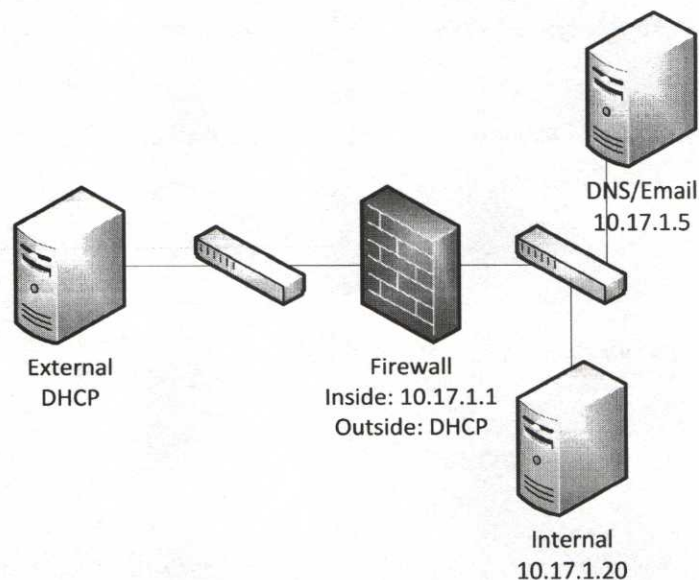
Are the unused switch ports locked down?

Do you have any observations about the VLAN configuration?

Please feel free to repeat these activities using the other two configuration files at your leisure! Hopefully, you now have a feel for why a tool like Nipper is so important. We hope that you can also see that having a general familiarity with how these devices function helps a great deal when trying to examine whether or not they are configured securely!

Part 2: Firewall Validation

In this section, we move from the analysis of the configuration files to actually validating how the firewall is behaving. To do so, let's start with a diagram to illustrate how this network is configured and where our various hosts are placed.



As you can see, it is a fairly simple network. In reality, when you perform a firewall validation, you should focus on the firewall and where you might place the scanners and sniffers.

In this case, you define the internal network to be the 10.17.1.0/24 network. Of course, this can be any network of any size at all. On the other hand, the outside of the network, where the External host sits, uses DHCP to assign addresses. This is somewhat particular to our situation because we are using VMWare and need to configure hosts in such a way that they would “just work” for anyone. For a firewall administrator working on this type of issue within the environment, there is no doubt that the firewall will have a static address. In that case, you would place your scanner/sniffer between the firewall and the border router.

One other item that we cannot illustrate in the lab using virtual machines is that we would prefer to connect the sniffer using either a span port on the switch or an actual network tap. To the right, you can see an example of a NetOptics gigabit tap. Taps come in all shapes and sizes, including fiber. Of course, the faster the tap, the more expensive it is. The tap pictured here would cost about \$700. You can obtain used taps for substantially less. For instance, a similar tap on eBay sells for \$79.99 as a “Buy it Now” item.



The need for a tap or span port exists because you need to see all of the inbound packets. It is common to find the firewall acting not only as a NAT, but also as port forwarding inbound traffic to specific hosts. If you cannot see all of the traffic, you might miss important indications that packets are leaking through the firewall!

Let's get started with the validation. To do so, you need to lay out a few simple firewall policies for this organization. After you have this documented, you can try to validate what *actually* happens!

Firewall Requirements:

Administration:

- Firewall administration is performed only on the internal interface over HTTP.

Inbound rules:

- Inbound connections to port 25 TCP will be forwarded to DNS/Email.
- Inbound connections on port 53 TCP and UDP will be forwarded to DNS/Email.
- Nothing else should be permitted inbound.

Outbound rules:

- Outbound traffic to ports 80 and 443 TCP are permitted to allow web traffic.
- Outbound traffic to port 25 TCP is permitted to allow outbound SMTP.
- Outbound traffic to port 53 UDP is permitted to allow DNS to function.
- No other traffic should be permitted outbound.

With these rules laid out, you can begin. Of course, you would have put together this list based on an information flow analysis and review of change control coupled with a review of the high-level policy documents and standards. To do the analysis, start by turning on the sniffer on the outside and firing packets outbound to it.

Please ensure that, at a minimum, the firewall, the Internal system, and the External systems are running. If possible, start the DNS server as well.

Once these systems are running, log into the External system and double-click the Wireshark icon. After this is done, please start the sniffer listening on interface eth0.

To start the sniffer, first open Wireshark. Once Wireshark is up and running, click the "Interface List" option.



With the capture interfaces displayed, select the “eth0” interface and click the “Start” button. When you do so, the Wireshark capture interface will open. Chances are that the window will remain blank, but it is possible that your computer might broadcast some packets every now and then.

Next, switch over to the Internal host. From the Internal host, let’s try to verify that outbound packets to ports 80 and 443 are permitted. To check this, take note of the IP address of your External host (the address is displayed on the desktop of the External system) and use HPing3 from a root shell to send TCP SYN packets to the External host on ports 80 and 443.

See whether you can answer the following questions:

Looking at the output from HPing, does it appear that the packets have successfully passed through the firewall?

How can you tell?

Can you be certain that this isn't the firewall rejecting the packets?

Please look at the output in the sniffer on the External system. You might need to scroll or you can set a filter to look for TCP packets. Did the packets from HPing arrive at the target host?

No.	Time	Source	Destination	Protocol	Length	Info
2	7.456500000	192.168.129.128	192.168.129.129	TCP	60	53544 >
3	7.456573000	192.168.129.129	192.168.129.128	TCP	54	http >
4	8.456630000	192.168.129.128	192.168.129.129	TCP	60	22493 >
5	8.456740000	192.168.129.129	192.168.129.128	TCP	54	http >
6	9.456545000	192.168.129.128	192.168.129.129	TCP	60	35933 >
7	9.456601000	192.168.129.129	192.168.129.128	TCP	54	http >
16	73.224115000	192.168.129.128	192.168.129.129	TCP	60	37828 >
17	73.224186000	192.168.129.129	192.168.129.128	TCP	54	https >
18	74.223938000	192.168.129.128	192.168.129.129	TCP	60	38461 >
19	74.223992000	192.168.129.129	192.168.129.128	TCP	54	https >
20	75.225754000	192.168.129.128	192.168.129.129	TCP	60	50160 >
21	75.225864000	192.168.129.129	192.168.129.128	TCP	54	https >

Pictured above you can see the output from a sniffer with the TCP filter applied. You can see the even numbered packets originate with the firewall. Why the firewall? Because it is acting as a NAT, concealing the internal 10.17.0.0/16 address space. You can see that these even-numbered packets go to the external host and that the External host then sends packets back from ports “http” (or 80) and “https” (or 443).

Using the sniffer, there is absolutely no question about what goes through the firewall! Next, let's see what happens when you send packets inbound. To do this, you're going to make a change to the configuration of the External system.

Close Wireshark. You can save the capture file if you'd like, but there is really no need.

After Wireshark has closed, open a root shell on the External system. When that shell is running, please execute the command, "route -n."

```
root@external:/home/auditor# route -n
Kernel IP routing table
Destination      Gateway          Genmask         Flags Metric Ref    Use Iface
0.0.0.0          192.168.129.2  0.0.0.0        UG    0      0      0 eth0
192.168.129.0    0.0.0.0         255.255.255.0  U     0      0      0 eth0
```

When you run the "route -n" command, you are asking the system to tell you how packets will be routed based on the networks that are available. The "-n" option asks it not to try to resolve those numbers back to host names.

Your output should look similar to what is above, except that the Gateway for Destination "0.0.0.0" will likely be different. The destination of all zeroes is what is known as the default route, or for a Cisco device, the gateway of last resort. In other words, unless there is another route with more specific information, any packet that we don't know what to do with will be given to the gateway at 192.168.129.2.

We want to add a route that tells the External system that if it ever wants to talk to a host on the 10.17.0.0/16 network, it should send that packet to the external address of the firewall!

Please switch to your firewall system and make a note of the IP address assigned to the WAN interface. If there is no address displayed or the address is listed as "Unknown," click the firewall and press the Return key to refresh the addresses.

```

Player | [Icons] | [Close] [Maximize] [Full Screen]
*** This is m0n0wall, version 1.8.1
    built on Wed Jan 15 13:32:38 CET 2014 for generic-pc.
    Copyright (C) 2002-2014 by Manuel Kasper. All rights reserved.
    Visit http://m0n0.ch/wall for updates.

LAN IP address: 10.17.1.1
WAN IP address: 192.168.129.128

Port configuration:

LAN   -> le1
WAN   -> le8

m0n0wall console setup
*****
1) Interfaces: assign network ports
2) Set up LAN IP address
3) Reset webGUI password
4) Reset to factory defaults
5) Reboot system
6) Ping host

Enter a number:

```

Notice that after pressing Enter on the firewall, the address “192.168.129.128” is displayed. You should expect your address to be different!

Now that you have this address, let’s add a route. *Using the address that you have just found, please run the following command from the root shell on your external virtual machine:*

```
route add -net 10.17.0.0/16 gw 192.168.129.128
```

Of course, you must replace the 192.168.129.128 with the address of *your* firewall. Running the “route -n” command should now produce output similar to this:

```

root@external:/home/auditor# route -n
Kernel IP routing table
Destination    Gateway         Genmask         Flags Metric Ref    Use Iface
0.0.0.0        192.168.129.2  0.0.0.0         UG    0     0     0 eth0
10.17.0.0      192.168.129.128 255.255.0.0     UG    0     0     0 eth0
192.168.129.0  0.0.0.0        _ 255.255.255.0  U    0     0     0 eth0

```

As you can see, we have successfully added the route.

Now that the route has been configured to force 10.17.0.0/16 traffic to the firewall, start Wireshark on the Internal host in the same way that you set it up on the External host. Again, you might see some extraneous packets being generated by your host. It is safe to ignore these.

When Wireshark is running on the inside, let’s try to send some inbound packets. Please send the following packets using HPing and answer the following questions:

- Packet to the firewall on port 53 UDP
- Packet to the firewall on port 53 TCP
- Packet to the firewall on port 22 TCP
- Packet to 10.17.1.20 on port 22 TCP
- Packet to 10.17.1.20 on any TCP or UDP port

Which of these packets appear to pass through the firewall to be received on the Internal host?

Did HPing correctly report responses in all cases?

No.	Time	Source	Destination	Protocol	Length	Info
5	1.975263000	192.168.129.129	10.17.1.5	DNS	60	[Malformed Packet]
6	16.736776000	192.168.129.129	10.17.1.5	TCP	60	h323gatestat > domain [
7	16.746710000	10.17.1.5	192.168.129.129	TCP	60	domain > h323gatestat [
8	16.750544000	192.168.129.129	10.17.1.5	TCP	60	h323gatestat > domain [
9	17.739056000	192.168.129.129	10.17.1.5	TCP	60	h323hostcall > domain [
10	17.739069000	10.17.1.5	192.168.129.129	TCP	60	domain > h323hostcall [
11	17.743277000	192.168.129.129	10.17.1.5	TCP	60	h323hostcall > domain [
41	27.275968000	192.168.129.129	10.17.1.20	TCP	60	pip > ssh [SYN] Seq=0 W
42	27.277151000	10.17.1.20	192.168.129.129	TCP	58	ssh > pip [SYN, ACK] Se
43	27.282152000	192.168.129.129	10.17.1.20	TCP	60	pip > ssh [RST] Seq=1 w
44	28.276090000	192.168.129.129	10.17.1.20	TCP	60	novation > ssh [SYN] Se
45	28.276172000	10.17.1.20	192.168.129.129	TCP	58	ssh > novation [SYN, AC
46	28.278033000	192.168.129.129	10.17.1.20	TCP	60	novation > ssh [RST] Se

In the previous, you can see a sample packet capture. Note that we have filtered the display so that only packets coming from the External scanner will be displayed. This allows you to see only the packets related to the tests that you just performed.

You can see that the inbound DNS packet was correctly sent to the system, whether it was sent over UDP or TCP. You can also see that the port 22 traffic was successful. You might have noticed some other items that were somewhat concerning, though! Let's look at some of the HPing output. *(Although we will look at it in HPing, this data is also visible in the Wireshark output. It is simply easier for us to read it in HPing to give the data context.)*

```

root@external:/home/auditor# hping3 -S -p 22 10.17.1.20
HPING 10.17.1.20 (eth0 10.17.1.20): S set, 40 headers + 0 data bytes
len=46 ip=10.17.1.20 ttl=63 DF id=0 sport=22 flags=SA seq=0 win=14600 rtt=5.6 ms
len=46 ip=10.17.1.20 ttl=63 DF id=0 sport=22 flags=SA seq=1 win=14600 rtt=2.1 ms
^C
--- 10.17.1.20 hping statistic ---
2 packets transmitted, 2 packets received, 0% packet loss
    
```

As you can see, when you sent packets from External directly to 10.17.1.20, they apparently passed through the firewall!! This is absolutely unexpected behavior that the administrator would need to analyze and correct!

Part 3 – Automated Firewall Validation

PLEASE NOTE! There is an alternate form of this exercise located in the appendix of this workbook beginning on page 228. If you are an auditor who will be called upon to perform a PCI validation of a firewall, then you will definitely want to review this extra lab. To understand the precise process that should be followed by administrators to validate your network controls whenever changes are made, you will also find this appendix very illuminating.

For our class, however, you are going to “fast forward” to the point that all of the “packet slinging” has been done and you are ready to actually perform the analysis. This will allow you to see how truly easy it is to do the validation, removing all of the manual effort that we just experienced using HPing and Wireshark.

Please bring the External Scanner to the foreground. If you are not already logged on, please log on now.

All of the files necessary to perform the actual analysis are already sitting on our External scanner. We’re simply going to leverage them to perform an analysis of our firewall very quickly!

After logging on, please start up a root shell by clicking the red “Root Terminal” icon on the desktop.

Again, day-to-day operations are never performed by logging in directly as root. In this case, you will work with files that have been created by the root user, however. The reason for this is that, as mentioned in class, a sniffer has been used to collect all of the packets that pass through the firewall in either direction. Running a sniffer on UNIX systems requires root-level privileges.¹

The files that you will work with are all located in the “PCI” directory that is contained in the “AuditcastsScripts” directory. ***Please change to the “AuditcastsScripts/PCI” directory.*** You can accomplish this by typing the following at the root command prompt:

```
cd AuditcastsScripts/PCI
```

¹ Alternatively, it is possible in many modern UNIX systems to grant users the ability to run a sniffer without that user becoming root. We have not taken this approach on our lab system because it is actually quite rare to find that administrators have done so.

Terminal (as superuser)

```
File Edit View Search Terminal Help
root@external:/home/auditor# cd AuditcastsScripts/PCI
root@external:/home/auditor/AuditcastsScripts/PCI# █
```

Now that you are in the proper directory, *type the “ls” command* to see the files that are in this folder on the External machine:

```
File Edit View Search Terminal Help
root@external:/home/auditor# cd AuditcastsScripts/PCI
root@external:/home/auditor/AuditcastsScripts/PCI# ls
cisp_analyze  cisp_sniffer    e_scan_settings  i_scan_settings
cisp_ciphers  cisp_vuln_scan  i_nmap_scan      outer_capture
cisp_scanner  e_nmap_scan     inner_capture     sample_pci_report.html
root@external:/home/auditor/AuditcastsScripts/PCI# █
```

In this directory, you can see the scripts that comprise the components of this free PCI/DSS validation tool (cisp_analyze, cisp_ciphers, cisp_scanner, cisp_sniffer, cisp_vuln_scan). If you’d like to actually use the scanner and sniffer, the lab in the appendix beginning on page 228 will walk you through this process.

Even if you do not want to work through the appendix lab, you might find it extremely useful when discussing the validation process with the network and firewall administrators. Feel free to show your administrators the directions found there or even to work through the lab themselves to see what we are looking for when it comes to validations.

To perform the analysis, please type the following command on the External host:

```
./cisp_analyze
```

This command will automatically discover which tests have actually been run, and then begin prompting you for information about what it has found. *The information for answering these questions can be found on page 49.*

There are two kinds of questions that you will be asked. First, are there ports that the *firewall* requires to be open on the inside or outside? In our case, because it is a NAT, we would expect any inbound services that are needed to be listening on the firewall itself.

The second type of question is about the *Infrastructure*. Are there ports that must pass *through* the firewall, though the firewall itself might not be listening on them?

SANS Advanced Systems Audit Workbook

```
Please enter the port number: 80
Are there more required ports? [y/N]
Added 1 permitted inbound ports on the inside of the firewall.
Are there ports that should be open on the outside of the firewall? [y/N]y
Please enter the port number: 25
Are there more required ports? [y/N]y
Please enter the port number: 53
Are there more required ports? [y/N]
Added 2 permitted inbound ports on the outside of the firewall.
Does your infrastructure require more outbound ports? [y/N]y
To enter a required port, you must first specify the
protocol type. Is this port a TCP or a UDP port? [T/u]t
tcp selected. Please enter the outbound port number: 25
tcp/25 added.
Does your infrastructure require more outbound ports? [y/N]y
To enter a required port, you must first specify the
protocol type. Is this port a TCP or a UDP port? [T/u]u
udp selected. Please enter the outbound port number: 53
udp/53 added.
Does your infrastructure require more outbound ports? [y/N]
reading from file inner_capture, link-type EN10MB (Ethernet)
reading from file inner_capture, link-type EN10MB (Ethernet)
reading from file outer_capture, link-type EN10MB (Ethernet)
reading from file outer_capture, link-type EN10MB (Ethernet)
Report completed. View 'pci_report.html' for the results.

root@external:/home/auditor/AuditcastsScripts/PCI# █
```

After answering these questions, we're ready to see the results. The report has been stored in a file called "pci_report.html." *To view the report, please type the following on the External system:*

```
firefox pci_report.html
```

You should see the Firefox web browser open with the report displayed!

Please examine this report and see whether there were any unexpected ports open on the firewall or through the firewall.

Exercise 4: Finding Wireless Clients

Time Required: Approximately 10 minutes

Purpose: Provide you with an easy-to-repeat process for identifying internal hosts connecting to external open access points. Illustrate exposures of using open wireless access networks.

Now that we've discussed secure deployment of wireless, there remains an important task. Although breaking into a network that makes of a Pre-Shared Key (PSK) is interesting, it is actually not particularly likely that our employees are breaking into encrypted networks to gain Internet access. What is very likely is that some employees are connecting to nearby unsecured access points from our office in an effort to bypass filtering and other security controls. This is a serious security matter because these systems might even be connected to the internal network simultaneously! As it turns out, there is a simple way to find hosts of this sort.

Scenario

For the last 45 minutes, a security engineer armed with a laptop and antenna concealed in a messenger bag has casually walked up and down all of the hallways and around the perimeter of the facility. During his walk, Airodump-NG was configured to capture all packets while hopping from channel to channel. He has just returned.

The security engineer has placed a copy of the packet capture file onto your computer in the “/home/auditor/AuditcastsScripts/wireless/wireless_capture.cap” file.

Please open a root command prompt on your External virtual machine. Once the command prompt is open, use the “cd” command to switch into the “AuditcastsScripts/wireless” directory. The root command shell is necessary because TCPDump is not in the path for “auditor.”

```
root@external:/home/auditor# cd AuditcastsScripts/Wireless/  
root@external:/home/auditor/AuditcastsScripts/Wireless# ls  
find_netbios netbios.pl wireless_capture.cap  
root@external:/home/auditor/AuditcastsScripts/Wireless# █
```

The capture file contains the raw data that was read by Airodump-NG. Although you can certainly open it in Wireshark, just display its contents quickly using TCPDump. *Please run the command “tcpdump -n -r wireless_capture.cap”.*

As TCPDump runs, the data just scrolls right off the screen. What we are looking for are the NetBIOS names that are being broadcast by all Windows hosts, and other hosts that are willing to participate on a Windows network. Finding these NetBIOS names using TCPDump is possible, but will require knowledge of how to create sniffer filters and the ability to understand the data that is being parsed.

It turns out that there is a much easier way to approach this problem. Because the NetBIOS names are simple encoded strings, we can just extract any strings that look like NetBIOS names and then decode them!

NetBIOS names are always 16 characters long. When encoded, these names will always be 32 characters long and those characters must fall in the range from uppercase "A" through uppercase "P". This makes finding these strings extremely easy! *Please enter the following command line:*

```
strings -l6 wireless_capture.cap
```

```
EFE0EDEMEBFGEFEGEPFCEFE0FDEJEDBM
70-35-10-73 AppleTV
70-35-10-73 AppleTV
ABACFPFPENFDECFCPEPFHFDEFFFPACAB
FEEJEOFEEBEEHEFEMCACACACACACAAA
EFE0EDEMEBFGEFCACACACACACACACABN
ABACFPFPENFDECFCPEPFHFDEFFFPACAB
FEEJEOFEEBEEHEFEMCACACACACACAAA
FEEJEOFEEBEEHEFEMCACACACACACAAA
FDENFDEDCACACACACACACACACACABN
\MAILSLOT\BROWSE
70-35-10-73 AppleTV
50-35-10-70.1 test
$!50-34-10-70.1 Nevada Time Capsule
70-35-60-63.1 Apple TV
70-35-10-73 AppleTV
50-35-10-70.1 test
$!50-34-10-70.1 Nevada Time Capsule
70-35-60-63.1 Apple TV
70-35-10-73 AppleTV
50-35-10-70.1 test
$!50-34-10-70.1 Nevada Time Capsule
70-35-60-63.1 Apple TV
root@external:/home/auditor/AuditcastsScripts/Wireless# █
```

As you can see, you are able to extract plaintext strings out of the packet capture without actually using a sniffer! In particular, please take special note of the lines that are all uppercase characters. These are precisely what we are looking for.

These NetBIOS names, however, are encoded. Decoding them requires a little bit of math, but we have a handy script that will do all of the work for us! *Please execute the following command line:*

```
./find_netbios wireless_capture.cap
```

When you execute this command, you should see what is shown below:

```

File Edit View Search Terminal Help
root@external:/home/auditor/AuditcastsScripts/Wireless# ./find_netbios wirele
ss_capture.cap
[hex]_MSBROWSE_[hex] <- Master Browser / Messenger Service
2012DC2 <- Domain Name / Workstation Service / IIS
ENCLAVE [hex] <- Master Browser
ENCLAVEFORENSIC[hex] <- Domain Controller / IIS
FENRIR <- Domain Name / Workstation Service / IIS
MACBOOKPRO-085F <- Domain Name / Workstation Service / IIS
SMSC [hex] <- Domain Master Browser
SMSC [hex] <- Master Browser
SMSC [hex] <- Browser Service Election Announcement
TINTADGEL <- Domain Name / Workstation Service / IIS
WORKGROUP [hex] <- Domain Master Browser
WORKGROUP [hex] <- Master Browser
WPAD <- Domain Name / Workstation Service / IIS
root@external:/home/auditor/AuditcastsScripts/Wireless# █

```

Notice that you can see machine names and domain names! The domain-related information has the funny characters after them. These are showing you the hexadecimal value that follows the NetBIOS name. Each of these codes is then decoded to the right-hand side, showing you exactly what type of NetBIOS name you are finding. We can also see computer names. Computer names do *not* have the hexadecimal characters after them (it's actually there, but you can't see it because it's a null) and is interpreted as either a Domain Name, a Workstation Service announcement, or an IIS Server announcement.

Of what use is this data? Imagine that our domain is ENCLAVEFORENSIC. Because we can see this name being broadcast in cleartext over some wireless link, we know for sure that at least one of the computers is communicating on that network. Every 30 seconds, this Windows host is announcing itself and its domain. A network engineer can now take this data into a sniffer like Wireshark and identify several things:

- Which wireless network was this seen on?
- What is the name of the machine probing for our domain?

Armed with this information, we can then visit the Domain Administrator who can identify which user logs into the domain from that workstation.

Exercise 5: DNS and SMTP

Time Required: Approximately 20 minutes

Purpose: Examine misconfigured SMTP and DNS servers. Provide repeatable remote testing processes for these systems that do not require the auditor to learn every possible configuration format for the large variety of servers available.

Requirements: This lab requires the External system, Firewall system, and DNS virtual machines. The Internal system is not required and can be shut down.

Part 1: DNS

The first tests will examine several interesting tests that should be run against our DNS servers in addition to illustrating why a Split DNS arrangement is so important. *To get started, please make sure that the External, DNS, and Firewall systems are all running.*

Log in to the External system and open a root prompt.

For all of these exercises, you MUST determine the WAN address of your Firewall system. Remember that this address is displayed on the console of the Firewall system. For the purposes of the instructions, we will use the address 192.168.129.128. You must replace this with your firewall address in all of the following examples.

DNS Version:

DNS servers should never reveal which version of the DNS service they are running. Another DNS server will never ask this question. It is, however, wonderful information for an attacker.

To attempt to retrieve the DNS version, enter the following commands:

```
nslookup
server 192.168.129.128
set class=chaos
set type=txt
version.bind
```

Does the server respond to this query? If it does, which version of the service is running?

```
root@external:~# nslookup
> server 192.168.129.128
Default server: 192.168.129.128
Address: 192.168.129.128#53
> set class=chaos
> set type=txt
> version.bind
Server:      192.168.129.128
Address:    192.168.129.128#53

version.bind  text = "9.8.1-P1"
> █
```

Exit out of nslookup to continue.

As you can see, the server reports that it is running version 9.8.1-P1. *Please search in google for “bind 9.8.1-P1 vulnerability.” Does it appear that this version has vulnerabilities?*

Zone Transfer:

Recall that a zone transfer allows a remote system to request a complete copy of a set of DNS records from a server. This feature is intended to be used only between DNS peers within the same zone, allowing the DNS information to be synchronized. Let's see how we can attempt a zone transfer manually. *Our sample DNS server is authoritative for the domain “Target.com.”*

Please run the following commands on the External system:

```
dig axfr @192.168.129.128 target.com
```

Was the command successful? If there are results, why would they be bad?

SANS Advanced Systems Audit Workbook

```
root@external:~# dig axfr @192.168.129.128 target.com

; <<>> DiG 9.8.4-rpz2+rl005.12-P1 <<>> axfr @192.168.129.128 target.com
; (1 server found)
;; global options: +cmd
target.com.                604800  IN      SOA     target.com. root.target.com. 1 6
04800 86400 2419200 604800
target.com.                604800  IN      NS      ns.target.com.
b2b.target.com.           604800  IN      A       170.129.50.10
b2b-dev.target.com.       604800  IN      A       170.129.50.11
firewall.target.com.      604800  IN      A       10.17.1.1
mail.target.com.          604800  IN      A       170.129.50.5
ns.target.com.             604800  IN      A       10.17.1.5
ns2.target.com.           604800  IN      A       170.129.50.3
oracle.target.com.        604800  IN      A       10.17.1.122
proxy.target.com.         604800  IN      A       170.129.50.25
research.target.com.      604800  IN      A       10.17.1.52
scanner.target.com.       604800  IN      A       10.17.1.20
dc.windows.target.com.    604800  IN      A       10.17.1.27
```

As you can see, the Dig command was successful. It turns out that even though NSlookup was *the* standard tool for interacting with DNS servers for many years, some modern versions do not actually implement all of the features. Dig, a newer tool, does allow you to test for this issue. Obviously, this is not good!!

Recursive Queries:

We would also like to verify that the DNS server will resolve only addresses for its own domain when a query comes in from the Internet. Let's try this out with the sample server.

Please execute the following commands on the External system:

```
nslookup
server 192.168.129.128
www.target.com
www.google.com
exit
```

Was the server able to resolve www.target.com?

Was the server able to resolve www.google.com?

If the server will resolve www.google.com, why is this bad?

SANS Advanced Systems Audit Workbook

```
root@external:~# nslookup
> server 192.168.129.128
Default server: 192.168.129.128
Address: 192.168.129.128#53
> www.target.com
Server:      192.168.129.128
Address:     192.168.129.128#53

Name:   www.target.com
Address: 170.129.50.2
> www.google.com
Server:      192.168.129.128
Address:     192.168.129.128#53

** server can't find www.google.com.localdomain: REFUSED
```

As you can see, the server does successfully resolve `www.target.com`, as it should. We also get our first piece of good news! The DNS server refused to resolve Google for us!

Split DNS:

We would also like to find that a split DNS arrangement is in use. To test for this, we will run two separate tests. *First, use the CD command to switch into the “AuditcastsScripts/DNS” directory on the External host.*

This directory has a handy Ruby script called “Mapper.rb.” To use the script, you send it a starting address, an ending address, and a DNS server to query.

Please use the mapper.rb script to query your DNS server for all of the hosts between address 170.129.50.0 and 170.129.50.255. The command line would look like this:

```
./mapper.rb 170.129.50.0 170.129.50.255 192.168.129.128
```

Are there any interesting hosts visible?

In this particular case, you actually end up with a sort of strange result. There were no answers at all!!

What this means is that the DNS server apparently has no reverse lookup entries populated for the 170.129.50.0/24 network. This is definitely unusual, but we wanted you to see what it would look like when there are no answers. Let's try something else. *Please enter the following command line:*

```
./mapper.rb 10.17.0.0 10.17.2.255 192.168.129.128
```

Are you able to see any hosts now?

Based on this answer, would you say that a Split DNS arrangement is in use?

SANS Advanced Systems Audit Workbook

```
root@external:/home/auditor/AuditcastsScripts/DNS# ./mapper.rb 10.17.0.0 10.17.2
.255 192.168.129.128
Starting at 10.17.0.0, counting up to 10.17.2.255
10.17.1.1:          firewall.target.com
10.17.1.5:          ns.target.com
10.17.1.20:         scanner.target.com
10.17.1.27:         dc.windows.target.com
10.17.1.52:         research.target.com
10.17.1.122:        oracle.target.com
root@external:/home/auditor/AuditcastsScripts/DNS# █
```

Pictured above is the output from this scan. Notice that you can see six different internal hosts whose information is both stored in and available from the public DNS server! Clearly, Split DNS is not in use. This is extremely serious because it can allow an attacker to map out where important systems are without ever running a network scanning tool. Worse, he can begin creating a mental map of the layout of our network, allowing him to quickly target the most interesting systems.

Part 2: SMTP

Let's turn our attention to SMTP configurations. Once again, you will make use of the DNS server in the lab environment.

Is there anything wrong with having our SMTP server housed on the DNS server?

Clearly the answer is "Yes!" We have violated the principle of separation of duties. We have also complicated the system. Now a vulnerability in our mail service leads to an immediate compromise of our DNS and vice versa.

Version/Headers:

To begin, let's see what the mail server is returning as header information. Just as with any other service, verbose headers can easily be used to attack the system. To do the testing, you will use a tool called "netcat." Netcat is not telnet, but you can certainly think of it that way. It allows you to establish a raw network connection to a remote system on any arbitrary port. *From the External system, please enter the following command:*

```
nc 192.168.129.128 25
```

When you enter this command, you will have to wait 20 to 30 seconds before you see any output (unless you have the IP address wrong... then you'll receive a connection failure!). What do you suspect the system is doing while you are waiting?

Mail servers are typically configured to perform a reverse lookup on any host that connects to them. This allows them to verify whether or not the remote host is lying about whom it is and can be used as a basis to refuse inbound email. *When the host does respond, what type of server is it, what platform is it on, and what version is it running?*

```
root@external:/home/auditor/AuditcastsScripts/DNS# nc 192.168.129.128 25
220 dns ESMTP Postfix (Ubuntu)
```

When the host does send its banner, you can see some information that you shouldn't see, but the information is limited. It is not ideal that the server is announcing that it is running Postfix, one of the many mail services that are available. It is also not ideal that it is telling you that it is running on an Ubuntu system. Still, it at least does not tell you the exact version of the software it is running.

VERFY and EXPN:

Next let's check to see whether VRFY or EXPN are enabled. Remember that these can be used to gather usernames and valid email addresses, possibly even determining the full names of users. *We will assume that you are already connected to the mail server from*

the previous portion of this lab. If you are not, please reconnect as you did in the previous step.

To test to see whether a user account exists with VRFY, you enter the command as follows:

VRFY username

Please check to see whether VRFY is supported by verifying the following usernames:

- auditor
- root
- dhoelzer

Is VRFY enabled? If it is, how can this be used by a phisher or spammer?

```
root@external:/home/auditor/AuditcastsScripts/DNS# nc 192.168.129.128 25
220 dns ESMTPE Postfix (Ubuntu)
VRFY auditor
550 5.1.1 <auditor>: Recipient address rejected: User unknown in local recipient
  table
VRFY root
252 2.0.0 root
VRFY dhoelzer
252 2.0.0 dhoelzer
```

Next let's try the EXPN command. *Please use the try the EXPN command in precisely the same way with the same users. Is this command enabled? If it is, how can this be used by a phisher or spammer?*

```
EXPN dhoelzer
502 5.5.2 Error: command not recognized
```

Although the VRFY command is enabled, you can see that EXPN is not. Good news!

Open Relay:

For the final test, let's see whether we can convince the server to send mail from an Internet user to another Internet user. We're also about to teach you how to forge email!

Mail servers have several commands that must be run in a particular order to send an email. First, you must say "Hello." Next, you have to say from whom the email is coming. Finally, you have to say where the email is going. *Using Netcat to connect to port 25 on the mail server, please try to send an email using the following process:*

```
nc 192.168.129.128 25
```

SANS Advanced Systems Audit Workbook

```
HELO spammer.com  
MAIL FROM: spammer@test.com  
RCPT TO: relay@there.com
```

If the server will accept mail as an open relay, you will now be prompted to send your message. *Is this server operating as an open relay?*

```
root@external:/home/auditor/AuditcastsScripts/DNS# nc 192.168.129.128 25  
220 dns ESMTP Postfix (Ubuntu)  
HELO spammer.com  
250 dns  
MAIL FROM: spammer@test.com  
250 2.1.0 Ok  
RCPT TO: relay@there.com  
554 5.7.1 <relay@there.com>: Relay access denied  
421 4.4.2 dns Error: timeout exceeded
```

You should investigate. For example, does the environment allow users to retrieve their email remotely? If so, how are credentials protected? How is the email content protected? What sort of authentication is used? In fact, these same questions would apply to an Internal email solution.

You also want to evaluate any anti-malware scanning or other features that serve to protect our endpoints. An additional question would investigate whether or not you should subscribe to any anti-spam services with the mail server.

Exercise 6: Network Discovery and Population Management

Time Required: Approximately 20 minutes

Purpose: Provide hands-on experience using a network scanner (Nmap). Introduce tools and systems for the effective long-term management of network scan data. Provide scenarios that illustrate the value of such systems.

Requirements: This lab uses the External and Firewall systems. Please ensure that these virtual machines are installed and running in order to complete this lab.

Nmap is probably the best-known network scanning tool in the world today. Although it has an enormous number of features, and we encourage you to experiment with them, we will look only at a subset that are particularly useful for building a network population management system.

Let's start by running a few basic scans. Remember that, as discussed in class, Nmap will first see whether the target host is up and, if it is, it will then port scan the host. Let's try to scan the firewall that we've been working with.

Open a root command prompt on your External system. Please begin by running the following command:

```
nmap -h | more
```

```
Nmap 6.00 ( http://nmap.org )
Usage: nmap [Scan Type(s)] [Options] {target specification}
TARGET SPECIFICATION:
  Can pass hostnames, IP addresses, networks, etc.
  Ex: scanme.nmap.org, microsoft.com/24, 192.168.0.1; 10.0.0-255.1-254
  -iL <inputfilename>: Input from list of hosts/networks
  -iR <num hosts>: Choose random targets
  --exclude <host1[,host2][,host3],...>: Exclude hosts/networks
  --excludefile <exclude_file>: Exclude list from file
HOST DISCOVERY:
  -sL: List Scan - simply list targets to scan
  -sn: Ping Scan - disable port scan
  -Pn: Treat all hosts as online -- skip host discovery
  -PS/PA/PU/PY[portlist]: TCP SYN/ACK, UDP or SCTP discovery to given ports
  -PE/PP/PM: ICMP echo, timestamp, and netmask request discovery probes
  -PO[protocol list]: IP Protocol Ping
  -n/-R: Never do DNS resolution/Always resolve [default: sometimes]
  --dns-servers <serv1[,serv2],...>: Specify custom DNS servers
  --system-dns: Use OS's DNS resolver
  --traceroute: Trace hop path to each host
SCAN TECHNIQUES:
  -sS/sT/sA/sW/sM: TCP SYN/Connect()/ACK/Window/Maimon scans
  -sU: UDP Scan
--More--
```

SANS Advanced Systems Audit Workbook

As mentioned earlier, Nmap has lots of features. Looking over the basic help, you can see that to run a scan, you simply need to execute Nmap and provide a target host address.

Please run Nmap against the WAN interface of your firewall system.

```
nmap 192.168.129.128
```

Which ports are open? Which services appear to be running?

```
root@external:/home/auditor# nmap 192.168.129.128

Starting Nmap 6.00 ( http://nmap.org ) at 2014-02-13 09:01 EST
Nmap scan report for 192.168.129.128
Host is up (0.0035s latency).
Not shown: 997 filtered ports
PORT      STATE SERVICE
25/tcp    open  smtp
53/tcp    open  domain
80/tcp    open  http
MAC Address: 00:0C:29:D1:15:AE (VMware)

Nmap done: 1 IP address (1 host up) scanned in 4.79 seconds
root@external:/home/auditor# █
```

When the scan executes, you will find that three ports appear to be open on the remote host. However, Nmap will actually perform its scan by scanning a set of common ports unless we tell it otherwise. ***Please use the help feature to find the option to specify target ports. When you have found this option, please rerun the scan, scanning for all ports from 1 through 65535.*** (Please note that this scan will take more than two minutes.)

Are any additional ports listening that were not visible before?

```
root@external:/home/auditor# nmap -p 1-65535 192.168.129.128

Starting Nmap 6.00 ( http://nmap.org ) at 2014-02-13 09:03 EST
Nmap scan report for 192.168.129.128
Host is up (0.0071s latency).
Not shown: 65532 filtered ports
PORT      STATE SERVICE
25/tcp    open  smtp
53/tcp    open  domain
80/tcp    open  http
MAC Address: 00:0C:29:D1:15:AE (VMware)

Nmap done: 1 IP address (1 host up) scanned in 130.44 seconds
```

You can see that the same ports are reported. This is good news from the point of view of firewall validation, but it certainly did take more time. *Let's try that again, but this time scan your External system.*

Does the scan run faster?

Why?

Which ports are open?

```
root@external:/home/auditor# nmap -p 1-65535 192.168.129.129
Starting Nmap 6.00 ( http://nmap.org ) at 2014-02-13 09:08 EST
Nmap scan report for 192.168.129.129
Host is up (0.00032s latency).
Not shown: 65532 closed ports
PORT      STATE SERVICE
22/tcp    open  ssh
111/tcp   open  rpcbind
53658/tcp open  unknown

Nmap done: 1 IP address (1 host up) scanned in 1.83 seconds
```

Wow, that was fast! Scanning the firewall took more than 2 minutes; scanning the own host took less than 2 seconds!!! Why is the firewall scan slower? Is it just because you have to pass over the network to scan it? No.

The reason that the firewall scan is so much slower is that the firewall is extremely unresponsive. Because it does not answer most probes, Nmap actually sends lots of extra probes just to make sure that it properly understands whether the ports are closed or whether the packets are being listed.

Using the help feature, please identify the option that performs OS Fingerprinting. Please rerun the scan against the Firewall and the External system, this time identifying the OS in use on each system.

Which OS does Nmap report for the Firewall?

Which OS does Nmap report for the External host?

SANS Advanced Systems Audit Workbook

```
root@external:/home/auditor# nmap -O 192.168.129.128-129

Starting Nmap 6.00 ( http://nmap.org ) at 2014-02-13 09:12 EST
Nmap scan report for 192.168.129.128
Host is up (0.0038s latency).
Not shown: 997 filtered ports
PORT      STATE SERVICE
25/tcp    open  smtp
53/tcp    open  domain
80/tcp    open  http
MAC Address: 00:0C:29:D1:15:AE (VMware)
Warning: OSScan results may be unreliable because we could not find at least 1 o
pen and 1 closed port
Device type: general purpose
Running (JUST GUESSING): Linux 3.X|2.6.X (91%)
OS CPE: cpe:/o:linux:kernel:3 cpe:/o:linux:kernel:2.6
Aggressive OS guesses: Linux 3.0 (91%), Linux 2.6.32 - 2.6.38 (87%), Linux 2.6.3
8 - 3.2 (85%)
No exact OS matches for host (test conditions non-ideal).
Network Distance: 1 hop

Nmap scan report for 192.168.129.129
Host is up (0.000048s latency).
Not shown: 998 closed ports
PORT      STATE SERVICE
22/tcp    open  ssh
111/tcp   open  rpcbind
No exact OS matches for host (If you know what OS is running on it, see http://n
map.org/submit/ ).
```

If you compare the command line used in the screenshot with the command line that you used, you might see a difference. Notice that we scanned both systems at the same time! It turns out that Nmap is very, very flexible when defining targets. For example, you can do something like this if you'd like:

```
Nmap 192.168.128-130.1-4,10-20,50,100
```

This allows you to surgically exclude or include target hosts. This is very handy when you are developing a scanning process and have hosts that are sensitive. It's also very useful if your networking team uses standards for host numbering on your subnets. For example, in most networks, the gateway that leads toward the core of the network will be at the .1 address. If you know this, you can quickly scan the entire IP space and identify which networks are alive by checking for .1 hosts. For example, imagine that the network address space was 128.226.0.0/16. To discover which networks are in use (again, supposing that we have the standard just mentioned), you could simply run:

```
Nmap -sP 128.226.*.1
```

This will scan all of the subnets to identify the .1 host. *What does the “-sP” option do?*

The “-s” option allows you to control the type of scan that you want to run. For example, an extremely valuable scan is the IP protocol scan. *Please locate the option for running*

a protocol scan. When you have found this option, run this scan against the External host and the Firewall.

Which protocols are available on the External host?

Which protocols are available on the Firewall?

```

root@external:/home/auditor# nmap -sO 192.168.129.128-129

Starting Nmap 6.00 ( http://nmap.org ) at 2014-02-13 09:20 EST
Nmap scan report for 192.168.129.128
Host is up (0.0016s latency).
All 256 scanned ports on 192.168.129.128 are open|filtered
MAC Address: 00:0C:29:D1:15:AE (VMware)

Nmap scan report for 192.168.129.129
Host is up (0.0000090s latency).
Not shown: 249 closed protocols
PROTOCOL STATE      SERVICE
1          open          icmp
2          open|filtered igmp
6          open          tcp
17         open          udp
103        open|filtered pim
136        open|filtered udplite
255        open|filtered unknown

Nmap done: 2 IP addresses (2 hosts up) scanned in 7.65 seconds
root@external:/home/auditor# █

```

So the “-sO” option will perform a protocol scan, but there’s something odd about our results! Notice that, apparently, the Firewall had no results! Actually, read the results more closely. Is it saying that there were no protocols? We know, for example, that TCP and UDP are definitely supported because we interacted with it previously.

What it actually reports is that all of the protocols are *filtered*. Because none of them answered, it doesn’t display any of them. On the other hand, the external host is definitely running three protocols. Four other protocols, however, were indeterminate, so it tells us that they might be filtered.

You might see similar output in port scans. If it knows the port is open, it will tell you. If it knows the port is closed, it can tell you. If it receives no answer at all, it has to assume that there is some kind of filtering happening.

Please identify the option that will perform Version scanning. When you locate this option, please use it to scan the Firewall and the External host.

Which version of the SSH service is running on the External host?

Which version of DNS is apparently running on the Firewall?

What kind of firewall does this appear to be?

Is there anything concerning about what you just discovered about the firewall?

```
root@external:/home/auditor# nmap -sV 192.168.129.128-129

Starting Nmap 6.00 ( http://nmap.org ) at 2014-02-13 09:28 EST
Nmap scan report for 192.168.129.128
Host is up (0.0024s latency).
Not shown: 997 filtered ports
PORT      STATE SERVICE VERSION
25/tcp    open  smtp      Postfix smtpd
53/tcp    open  domain    ISC BIND 9.8.1-P1
80/tcp    open  http      m0n0wall FreeBSD firewall web interface
MAC Address: 00:0C:29:D1:15:AE (VMware)
Service Info: Host: dns; OS: FreeBSD; Device: firewall; CPE: cpe:/o:freebsd:freebsd

Nmap scan report for 192.168.129.129
Host is up (0.0000080s latency).
Not shown: 998 closed ports
PORT      STATE SERVICE VERSION
22/tcp    open  ssh       OpenSSH 6.0p1 Debian 4 (protocol 2.0)
111/tcp   open  rpcbind
Service Info: OS: Linux; CPE: cpe:/o:linux:kernel

Service detection performed. Please report any incorrect results at http://nmap.org/submit/ .
Nmap done: 2 IP addresses (2 hosts up) scanned in 17.50 seconds
```

Using the “-sV” option, you can ask for version probing. What Nmap is actually doing here is called “Nudging.” It is simply sending a few carriage returns or similar to each open port and seeing what comes back.

Looking at the results, the port 80 result from the firewall is pretty concerning. Why so? If port 80 were open and reporting an Apache server or an IIS server, we would assume that we were looking at port forwarding. In this case, however, it is reporting Monowall! That is a firewall brand. It appears that the firewall has been misconfigured, allowing users to access the administration interface from the Internet!

Let’s change gears. Let’s see how to produce useful information from Nmap scans. ***Please find the option that will output the Nmap scan results as an XML file. When you find that option, please scan hosts 127.0.0.1 through 127.0.0.50 and store the results into a file named “first_scan.xml.”***

Nmap actually supports output to multiple different types of files. This can be useful for searching through and managing results. Let's use XML because there are some handy tools that will process this data into meaningful reports for us.

```
root@external:/home/auditor# nmap -oX first_scan.xml 127.0.0.1-50

Starting Nmap 6.00 ( http://nmap.org ) at 2014-02-13 09:35 EST
Nmap scan report for localhost (127.0.0.1)
Host is up (0.0000080s latency).
Not shown: 997 closed ports
PORT      STATE SERVICE
22/tcp    open  ssh
25/tcp    open  smtp
111/tcp   open  rpcbind

Nmap scan report for 127.0.0.2
Host is up (0.0000070s latency).
Not shown: 998 closed ports
PORT      STATE SERVICE
22/tcp    open  ssh
111/tcp   open  rpcbind

Nmap scan report for 127.0.0.3
Host is up (0.0000070s latency).
Not shown: 998 closed ports
PORT      STATE SERVICE
22/tcp    open  ssh
111/tcp   open  rpcbind
```

In this case, you are taking advantage of the fact that any address beginning with 127 will resolve to the local host. This allows you to simulate a network of 50 hosts very easily. When the scan completes, you will have an XML file of results.

Please execute the following command:

```
service rpcbind stop
```

```
root@external:/home/auditor# service rpcbind stop
[ ok ] Stopping rpcbind daemon....
root@external:/home/auditor# █
```

This will shut down the RPC service. This will force a change to the “network.” ***Please run the same scan again, this time storing the results into a file named “second_scan.xml.”***

Now that you have run these two scans, you can use some additional software to compare them. Let's start with the tool that's built into Nmap. ***Please execute the “ndiff-h” command.***

Ndiff is a basic tool for comparing to Nmap scans. ***Please execute the following command:***

SANS Advanced Systems Audit Workbook

```
ndiff first_scan.xml second_scan.xml | more
```

```
root@external:/home/auditor# ndiff first_scan.xml second_scan.xml | more
-Nmap 6.00 scan initiated Thu Feb 13 09:35:54 2014 as: nmap -oX first_scan.xml 1
27.0.0.1-50
+Nmap 6.00 scan initiated Thu Feb 13 09:39:21 2014 as: nmap -oX second_scan.xml
127.0.0.1-50
```

```
localhost (127.0.0.1):
-Not shown: 997 closed ports
+Not shown: 998 closed ports
PORT      STATE SERVICE VERSION
-111/tcp  open  rpcbind
```

```
127.0.0.10:
-Not shown: 998 closed ports
+Not shown: 999 closed ports
PORT      STATE SERVICE VERSION
-111/tcp  open  rpcbind
```

```
127.0.0.11:
-Not shown: 998 closed ports
+Not shown: 999 closed ports
PORT      STATE SERVICE VERSION
-111/tcp  open  rpcbind
```

Ndiff is now displaying for you changes that have been detected in the scans! Although this output is not terribly difficult to understand, it certainly isn't attractive. Items that have been added are preceded with a "+", items that have been removed are marked with a "-". Although you can understand it, you can do much better.

Please execute the following commands – please note that the second command is all one line but has been split over two within the text of the book:

```
yandiff --gen-stylesheet nmapstyle.xsl
```

```
yandiff --format xml --output-file nmap_report.xml --stylesheet
nmapstyle.xsl --baseline first_scan.xml --observed second_scan.xml
```

The first command generates a stylesheet that will be used for "styling" the report. The second command compares the first scan to the second and generates an XML report. ***To read the report, please execute "firefox nmap_report.xml" from the command line on your External host.***

New														
Missing														
Changed														
127.0.0.1 (localhost) - up														
<table border="1"> <thead> <tr> <th>Change Type</th> <th>Port</th> <th>Protocol</th> <th>Status</th> <th>Name</th> </tr> </thead> <tbody> <tr> <td>Missing</td> <td>111</td> <td>tcp</td> <td>open</td> <td>rpcbind</td> </tr> </tbody> </table>					Change Type	Port	Protocol	Status	Name	Missing	111	tcp	open	rpcbind
Change Type	Port	Protocol	Status	Name										
Missing	111	tcp	open	rpcbind										
127.0.0.2 - up														
<table border="1"> <thead> <tr> <th>Change Type</th> <th>Port</th> <th>Protocol</th> <th>Status</th> <th>Name</th> </tr> </thead> <tbody> <tr> <td>Missing</td> <td>111</td> <td>tcp</td> <td>open</td> <td>rpcbind</td> </tr> </tbody> </table>					Change Type	Port	Protocol	Status	Name	Missing	111	tcp	open	rpcbind
Change Type	Port	Protocol	Status	Name										
Missing	111	tcp	open	rpcbind										

If you peruse this report, you can see that it is *far* more user friendly than what NDiff creates. As a class, we'll discuss the answers to the following questions:

Who should be interested in seeing this report daily or weekly?

When would an auditor be interested?

What might an auditor ask to see in this kind of report?

How could this type of report be generated automatically within the enterprise?

Day 3

Exercise 1: Intro to HTML and HTTP

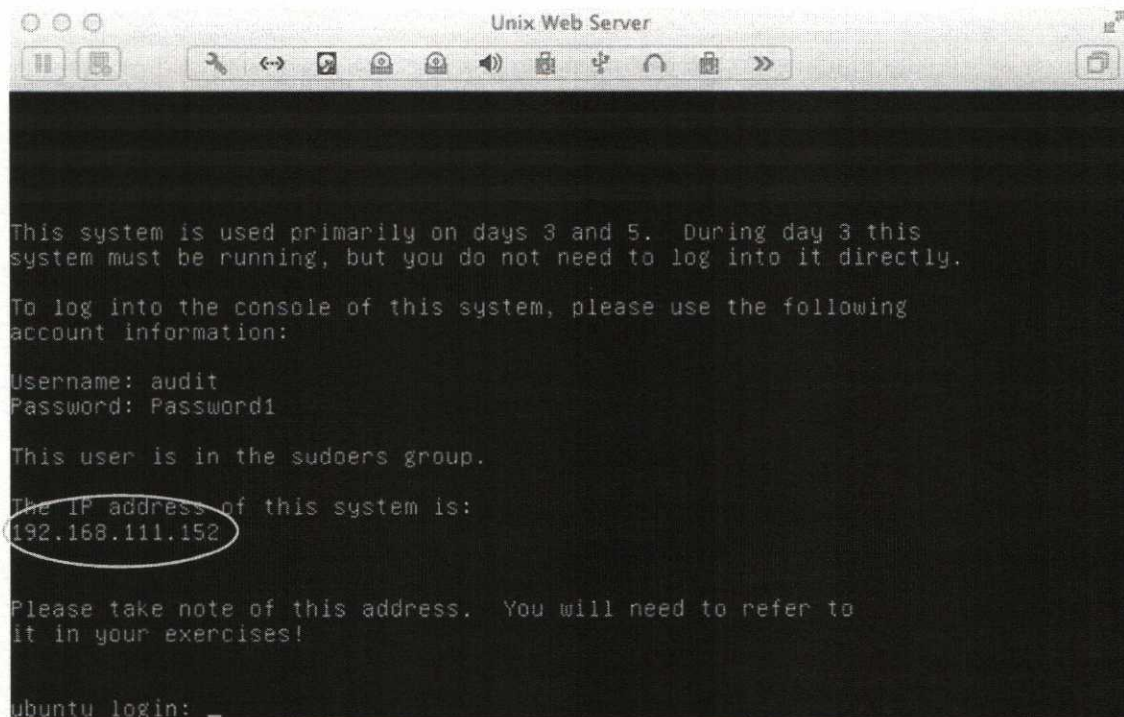
Time Required: Approximately 25 minutes

Purpose: Introduce HTML formatting and HTTP Protocol, allowing you to identify aspects of HTML and otherwise hidden information.

To begin this exercise, please open your web browser. During this course, the instructions for the exercises will be written with Internet Explorer in mind. Rest assured, this is not because Internet Explorer is the author's browser of choice; it's simply because it is nearly universally understood. All of the directions involving Internet Explorer can be accomplished using other browsers. It will be up to you to find the necessary menu options if he or she chooses to use a different browser.

You will also need one of the VMware systems running. Please start up the "Unix Web Server" system in VMware.

When the Unix Web Server system has completed the boot process, you will see a screen indicating local login information for the system in addition to the IP address that has been assigned.



```
Unix Web Server

This system is used primarily on days 3 and 5. During day 3 this
system must be running, but you do not need to log into it directly.

To log into the console of this system, please use the following
account information:

Username: audit
Password: Password1

This user is in the sudoers group.

The IP address of this system is:
192.168.111.152

Please take note of this address. You will need to refer to
it in your exercises!

ubuntu login: _
```

Make note of the address displayed in your virtual machine. You will need this address for all of the exercises for the remainder of today!

Using the address identified in the previous step, use your web browser to connect to that site. If the address was 172.16.81.128, you would direct your browser to “http://172.16.81.128.” Using this page and your browser, answer the following questions:

1. **Are there any “hidden messages” or errors that are evident?**
(Hint: Hidden messages are anything that you cannot see in the web browser, excluding tags that are present solely for formatting instructions.)
2. **How could you examine the page more closely?**
(Hint: Is there a feature in the browser that allows you to view the source code?)
3. **While examining the page using the alternative methods from question 2, are you able to find any additional information that is not readily apparent while looking at the page?**
4. **Can you find any “hidden” information?**
(Hint: Hidden messages are anything that you cannot see in the web browser, excluding tags that are present solely for formatting instructions.)
5. **If you found a hidden message, what type of HTML is it?**
6. **How could the information in our hidden message be useful to an attacker?**

It is not uncommon to discover hidden comments in web pages that give out confidential information or information about the inner workings of the web apps on the server. This is an area you would definitely want to have a look at if you were to audit web applications.

Hints:

If you view the page using the "View Source" option, you can search through and find several HTML comments. This is interesting but tedious. It is far more efficient to search using the "Edit->Search" or "Find" option, depending on which editor you are viewing the source code in. We would recommend that you search for tags like "<!--" and "<SCRIPT."

```

index[1] - Notepad
File Edit Format View Help

<body BGCOLOR="#000000" TEXT="#FFFFFF" LINK="#CCCCFF" ALINK="#FF0000" VLINK=
<p align="center"><font face="Verdana">n-log database query</font></p>
<!--
change the line below to your systems location of nlog-search.pl
//-->
<!-- There is a very serious bug in the OS Matching code that could potentially
allow a remote user to run arbitrary commands on the system. The CGI below
has the OS Match disabled, but a copy of the original was backed up. -->
<form action="/cgi-bin/nlog-search.pl" method="POST">
  <table border="0" width="100%" cellpadding="2">
    <tr>
      <td width="13%"><font face="Verdana" color="#FFFFFF"><small>Database</
      <td width="87%"><input type="text" name="database" size="40" value="sa
    </tr>
  </table>

```

Figure 1 - Third instance found searching for "<!--"

```

index[1] - Notepad
File Edit Format View Help

<tr>
  <td width="13%"><font face="Verdana" color="#FFFFFF"><small>OS Match</
  <td width="87%"><input type="text" name="os" size="20" value="*"></td>
</tr>
<tr>
  <td width="13%"><input type="submit" value="Search" name="submit"></td
  <td width="87%"></td>
</tr>
</table>
<input type="hidden" value="Hacker's Paradise" name="Source" >
</form>

<hr size="1" color="#C0C0C0">
<p align="center"><font face="Verdana">help with nlog-search.pl</font></p>
<p align="left">&nbsp;</p>
<p align="left"><small><strong><font face="Verdana">Database:</font></strong>

```

Figure 2 - First instance found searching for "hidden"

SANS Advanced Systems Audit Workbook

When examining the HTML comments, notice that one of them speaks of a very serious vulnerability in the code that is known. It also mentions that the code has been backed up. Where? The code doesn't say, but by looking at the name of the CGI program, perhaps we could poke around, trying to add "bak" to it or something of that sort, seeing whether perhaps it has been copied somewhere else on the server. Even if you can't run the code, if you can get a copy of the code, it could be used to find a vulnerability that is not known.

The last piece of hidden content on this page is found in a "Hidden" form field. Another excellent thing to search for is the key word "hidden." In this case, it reveals an element in the search form that you could not see before.

Part 2: Man-in-the-Middle Proxies

Purpose: To successfully use WebScarab and Burp Suite while introducing the basic proxy functionality of the tools. Establish basic familiarity with some of the functionality of the tools to facilitate later discussions and exercises.

Install JRE

To accomplish this exercise, you need to verify that a current version of Java is installed. If you already have Java installed, you can skip this step. If you do skip this step and then find that the tools will not function, you must return and follow these directions.

Installing Java requires you to connect to the Internet provided in class and connect to <http://www.java.com>.



The image shows a screenshot of the Java website homepage. At the top, there is a dark navigation bar with the Java logo on the left, a search bar in the center, and the words "Download" and "Help" on the right. Below the navigation bar, the main content area features the text "JAVA + YOU, DOWNLOAD TODAY!" in large, bold, sans-serif font. Underneath this text is a dark button with the text "Free Java Download" in white. Below the button, there are three small links: "What is Java?", "Do I have Java?", and "Need Help?". To the right of the text and button is a photograph of a man with glasses, wearing a white shirt, sitting on the floor and working on a laptop.

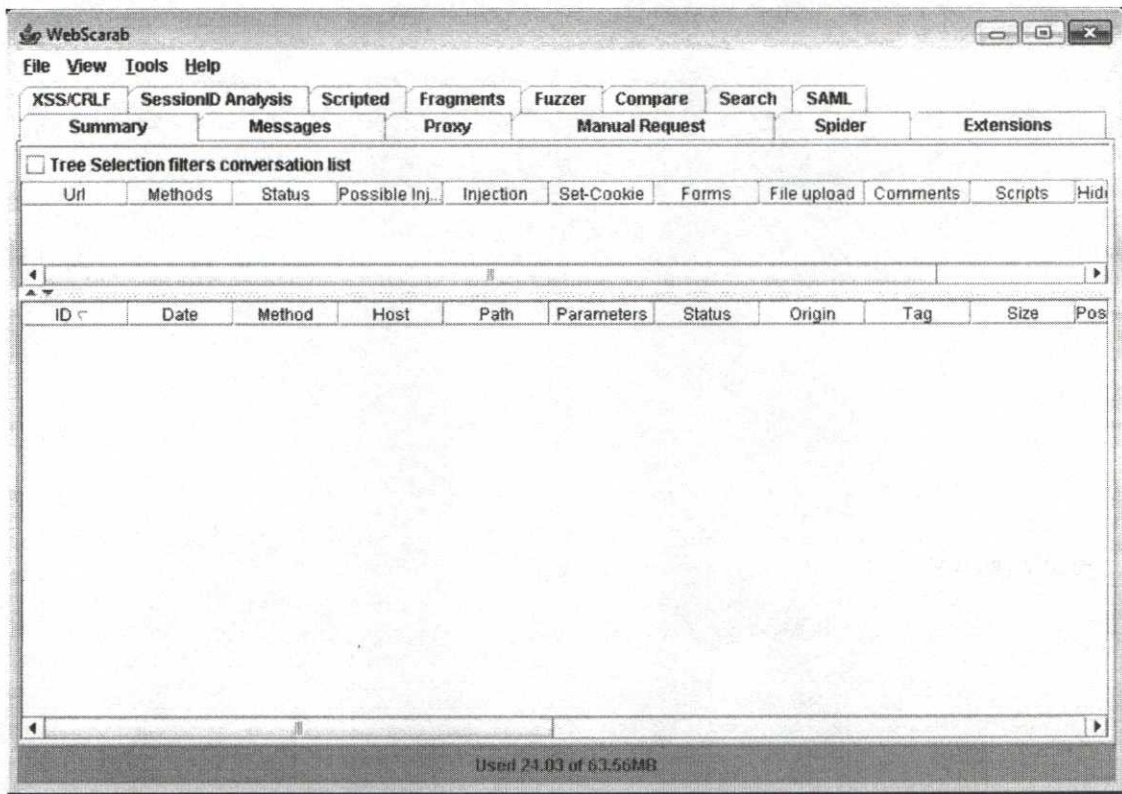
After connecting to <http://www.java.com>, please click the “Free Java Download” and follow the directions to install a working version of Java for your operating system.

WebScarab


To use WebScarab, simply execute the WebScarab JAR file within the “Web Application Testing” folder, which is found in the “Tools” folder that you copied from the USB.

Among other things, WebScarab will allow you to spy on and edit conversations with web servers by acting as a proxy server; in other words, you don’t connect directly to the web server, you connect to WebScarab and then WebScarab connects to the web server. This allows you to control what’s sent back and forth between the two.

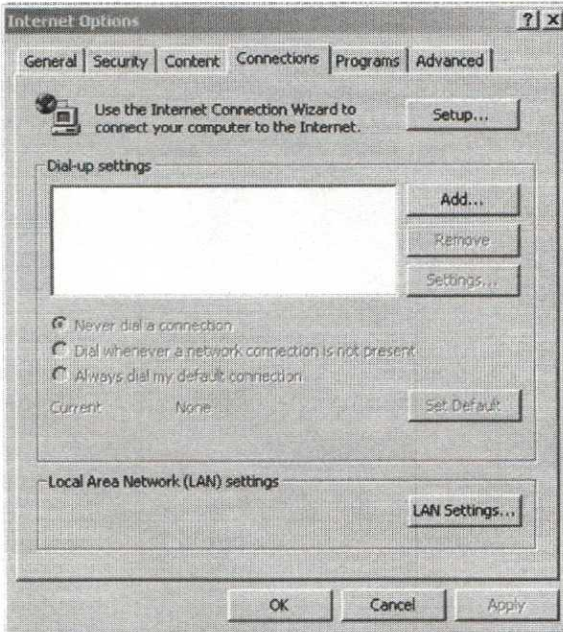
Please start up WebScarab now.



First, let's set up our browser to work with WebScarab. If we were to switch to the "Proxy" tab and the "Listeners" sub-tab, we would find that a listener has started on 127.0.0.1, port 8008. We need to configure our web browser to use this port for proxy connections.

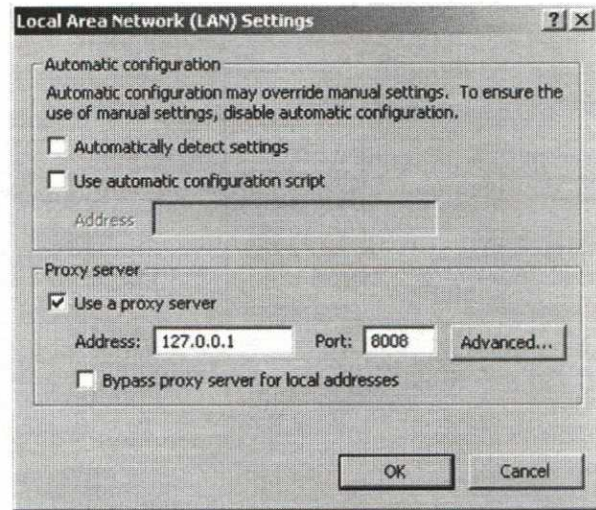


Locate the proxy settings for your web browser. If you are using Internet Explorer, you will find this setting under the "Tools" menu. Pull down this menu and select "Internet Options."



In the Internet Options dialog box, please locate and select the "Connections" tab along the top edge. Selecting the "Connections" tab will result in the dialog box to the left.

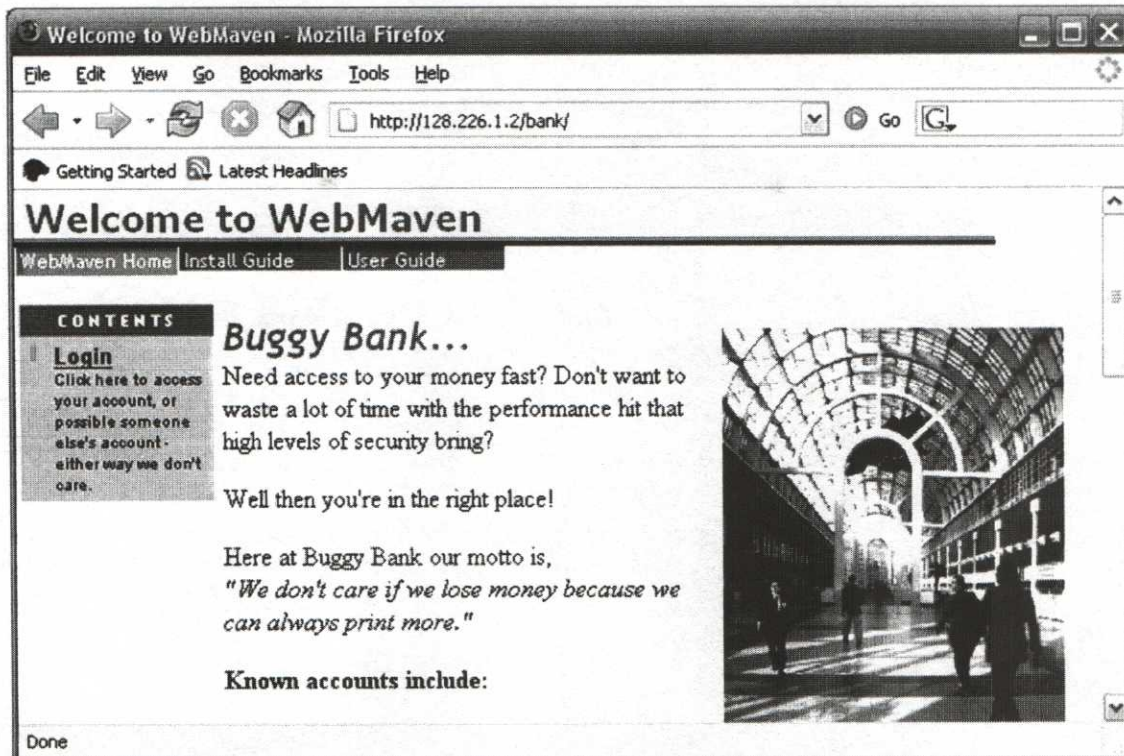
Please click the "LAN Settings" button in the "Connections" tab. This is where the actual proxy settings can be found. Please check the box for "Use a Proxy Server" and then fill in the blanks with host IP address 127.0.0.1 and port number 8008. Finally, click "OK" on each of the open dialog boxes so that we can test our configuration.



At this point, you should have configured the web browser to use WebScarab as a local proxy server when interacting with remote web servers. To test everything before proceeding, you will try to open a web page from the "Unix Web Server" system.

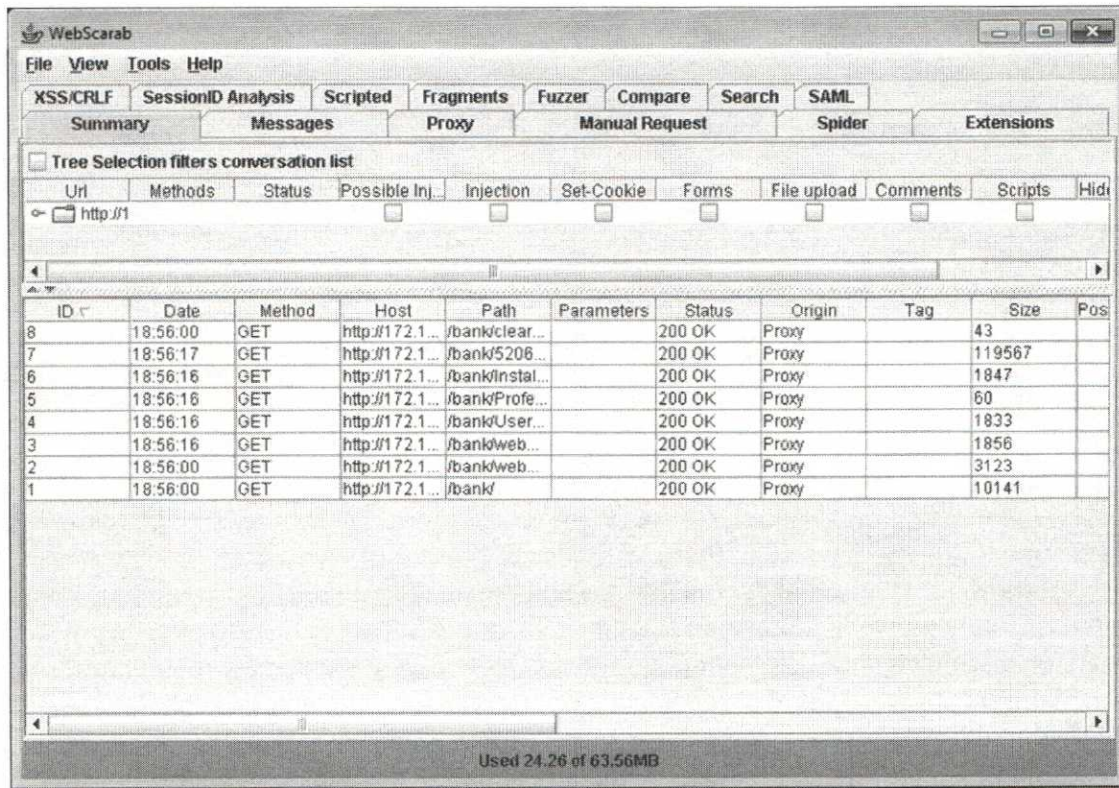
If it is not already running, please start up the "Unix Web Server" system.

After the "Unix Web Server" system starts up, please try to connect to the IP address of your virtual machine at the following URL: "http://xxx.xxx.xxx.xxx/bank/." If everything is working correctly, our web browser should display the following:



This on its own is not enough to guarantee that everything is correctly configured. We must also check the WebScarab interface to make sure that it is actually acting as a proxy server. Please bring WebScarab to the front. You should see something similar to this:

SANS Advanced Systems Audit Workbook

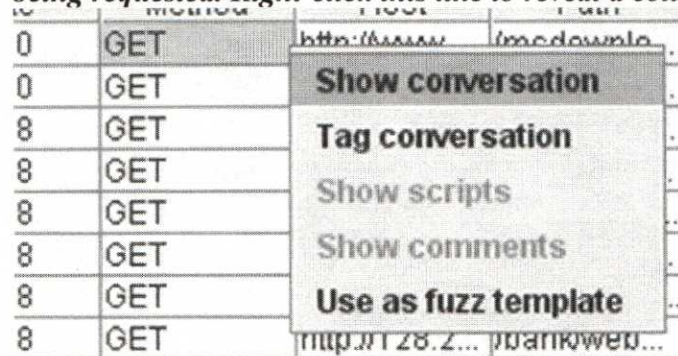


This screen shows the requests that have passed through WebScarab. Because you can see several pages were retrieved through the proxy, you can be sure that WebScarab and the web browser are properly configured.

By doing this, you told the web browser to direct all of its requests to the local machine (127.0.0.1) and port 8008. Also, if your browser asks whether it should accept a cookie from Buggy Bank, please say YES or you'll have a very hard time with the exercises.

One of the primary purposes of using a tool like this is to allow you to intercept and inspect the data being sent back and forth between the client and the server. The next steps will allow you to do this.

Please highlight one of the lines in the Summary view where the "/bank/" URL was being requested. Right-click this line to reveal a context menu:



Click the “Show Conversation” option.

What kind of information is revealed?

What kind of information is sent from the web browser to the web server?

What kind of information is sent back from the web server to the browser?

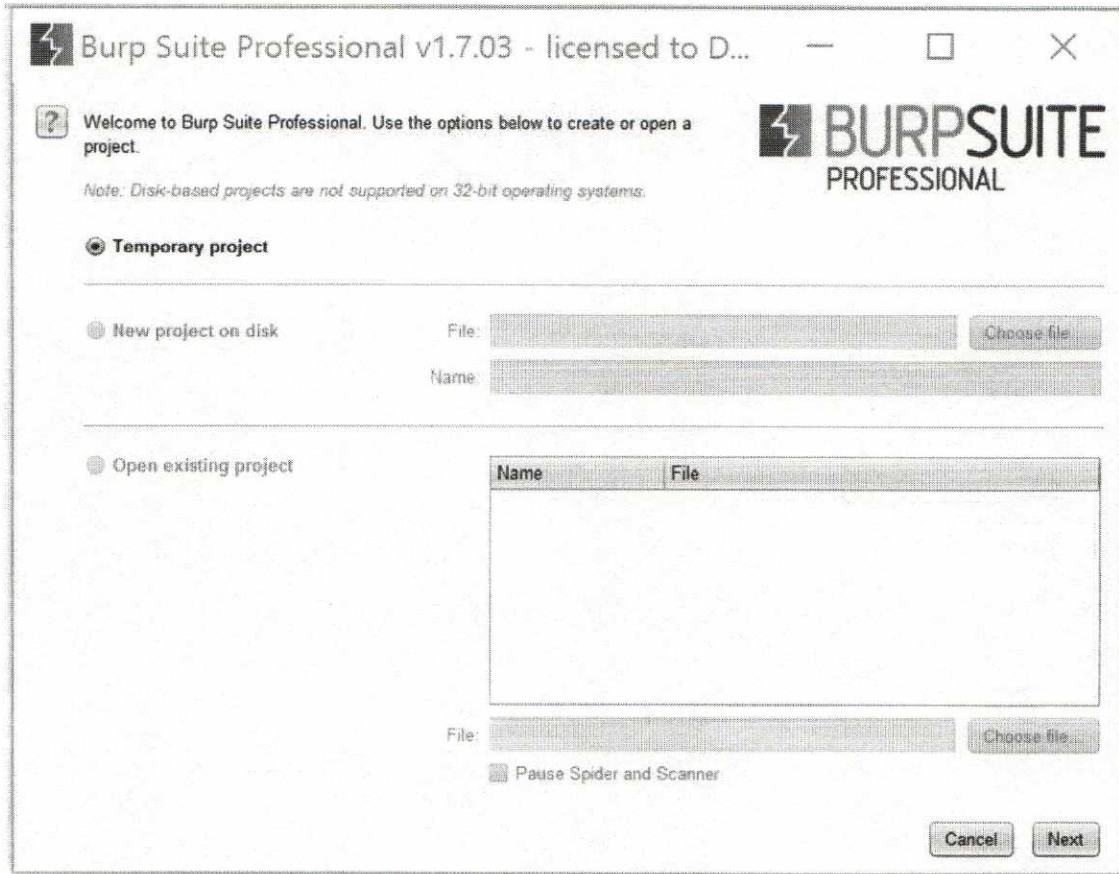
At this point, you can see that this tool can easily be used to retrospectively analyze the contents of requests and responses to a web application. We’re going to make good use of this feature as the day continues. This tool is also useful for actually testing for vulnerabilities too! To do this, you will want to do some “real-time” modifications on the web traffic as it passes by the proxy. For now, you are not going to do anything especially tricky. We just want to introduce the basics of the tool so that when we talk about more advanced topics, you’ll already have the toolkit to do whatever testing you can think of.

Burp Suite

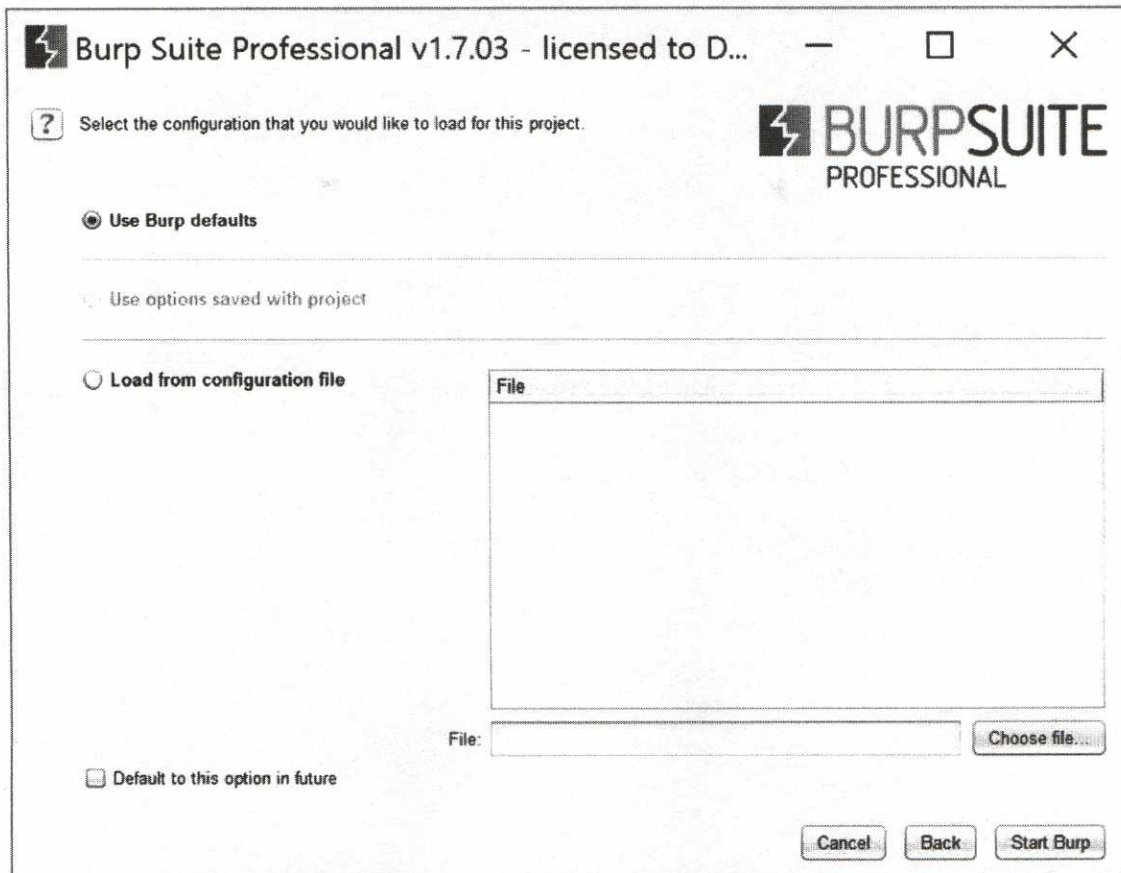
The Burp Suite tool is also in the “Tools” folder in the “Web Application Testing” folder where WebScarab was found. To use this tool, we again recommend that you copy the folder containing the Burp Suite to your local hard drive.

After Burp has been copied onto your system, please browse into the Burp folder and use the “burpsuite_free_v1.6.jar” icon to start this Java program.

The Burp logo will appear briefly if Java has been successfully installed. Immediately following that, Burp will prompt you to determine what type of project you would like to work on. *For our purposes, the ‘Temporary Project’ will be just fine.* The idea of a project is to allow you to organize all of the web requests, findings, notes and other items easily.

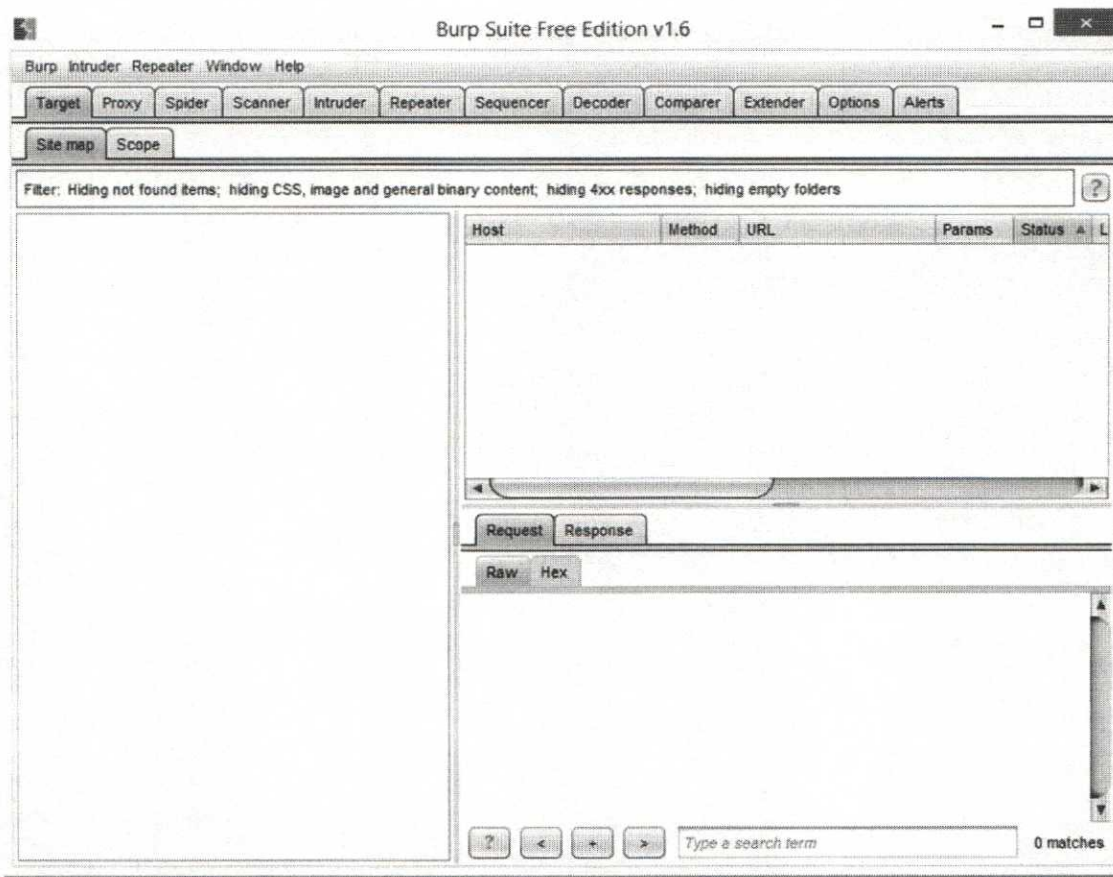


Pictured above is the project selection screen. Please leave Burp configured with the 'Temporary Project' default and select 'Next.'



Pictured above is the result of selecting a temporary project. At this point you have the option to load a different set of configuration defaults for your testing. ***Please leave Burp set to 'Use Burp defaults' and click on the 'Start Burp' button.***

SANS Advanced Systems Audit Workbook




When Burp starts up, it will start a proxy listening on port 8080. To make use of it, you will need to reconfigure the browser to point to 8080 rather than 8008, which is where WebScarab is running. ***Please change the proxy settings on the browser to use the Burp proxy at port 8080 now.***

There are two other options that can make this process easier to manage. First, there are some free "Proxy Switcher" tools available for the web browsers used on most Windows based systems. These will allow you to configure multiple proxies, and then switch between them with just a few clicks of the mouse.

Another option is to use what is known as "Proxy Chaining." What this means is that we can point one proxy at another proxy and thus chain them together. For instance, you could point your browser at WebScarab, and then point WebScarab at Burp. Now when you browse a web page, it will be available in both WebScarab and Burp. We are not going to use this configuration in the class because even though it sounds easier, it can actually become quite confusing, causing side effects as we're browsing pages.

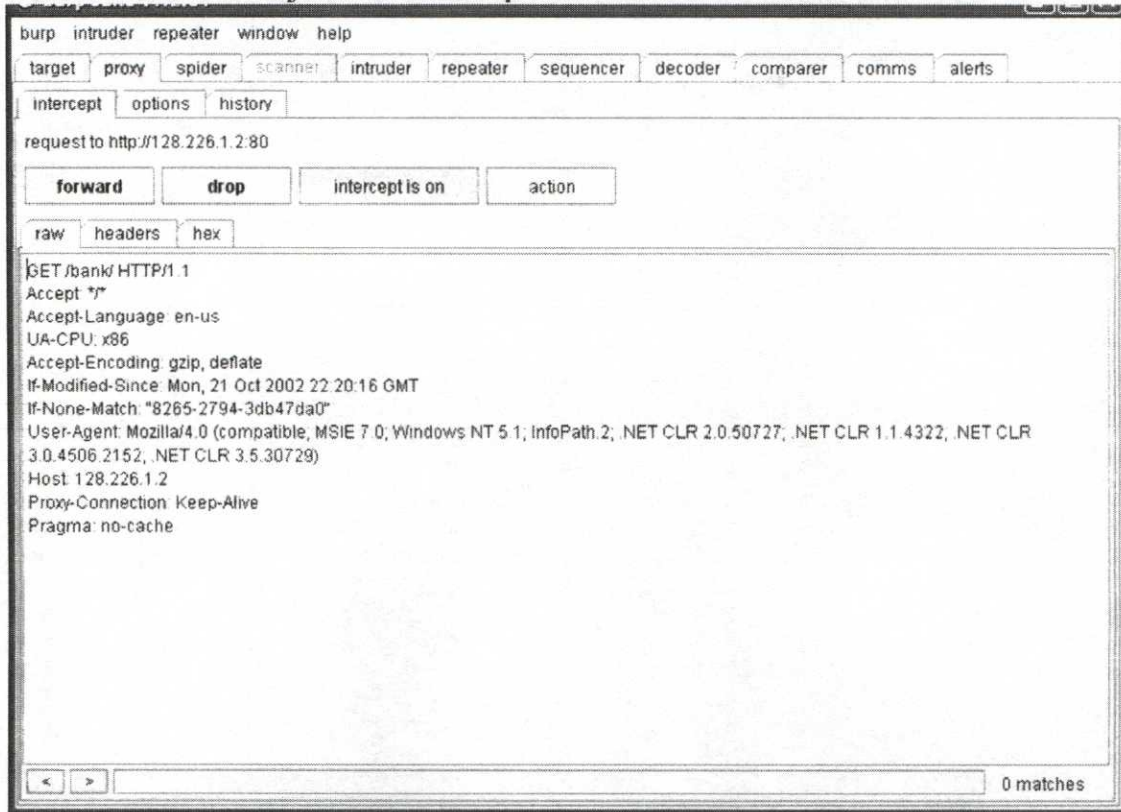
It is important to note that Burp starts up in "Intercept" mode. Please try to refresh the BuggyBank Homepage after setting Burp to be your current proxy.

When you attempt to reload the page, you should notice that the "Proxy" tab in Burp Suite turns red. You might also notice that the



page does not reload! The reason is that the page has been intercepted.

Please click the “Proxy” tab in the Burp Suite window.



You will notice that within the Burp Suite window under the Proxy tab, we can see the current request that has been intercepted. It is now available to be edited, dropped, or simply forwarded with no changes.

For now, please click the “Intercept is on” button to turn intercept off. You might also “forward” the current request to allow the refresh to complete.

As you work through the material today, WebScarab will be the primary tool, but please feel free to experiment with Burp as well. For a few of the exercises, we will work with both WebScarab and Burp to compare the functionality available in the two tools.

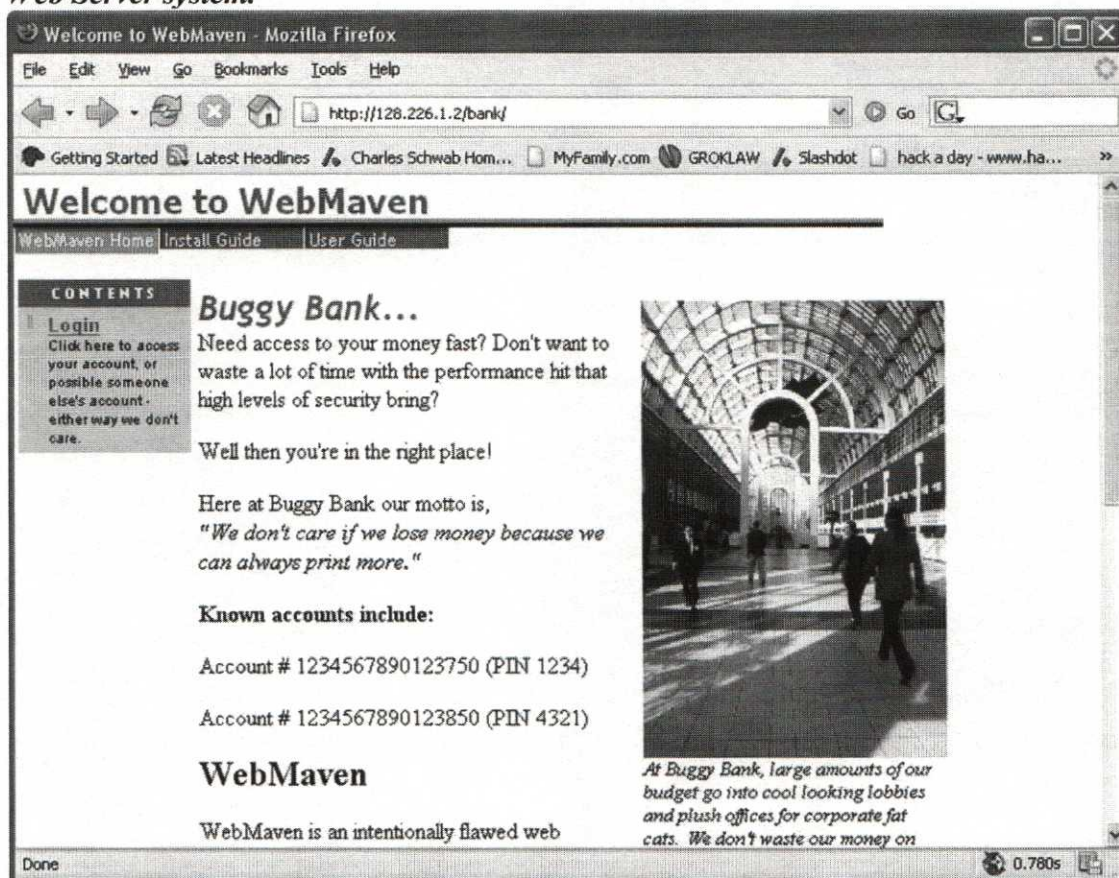
Exercise 2: Application Fuzzer

Purpose: To introduce the you to the concepts of automated fuzzing and demonstrate how to implement simple fuzzing tests using Burp.

For this exercise, you will need Burp running, your browser properly configured to use Burp as a proxy, and the Unix Web Server host up and running. With these prerequisites satisfied, let's get started!

As discussed during the lecture, fuzzing allows you to quickly and easily test parameters in a web application automatically, and then review the results. When we consider how many different combinations of tests are required to thoroughly exercise an application, this is a real time saver! Of course, this is also a technique that attackers use to automatically find potential errors in web applications. Remember, too, that this same technique (but using different tools) is used to exercise other types of applications as well.

To get started, please use your web browser to go to the "/bank/" URL on the Unix Web Server system.

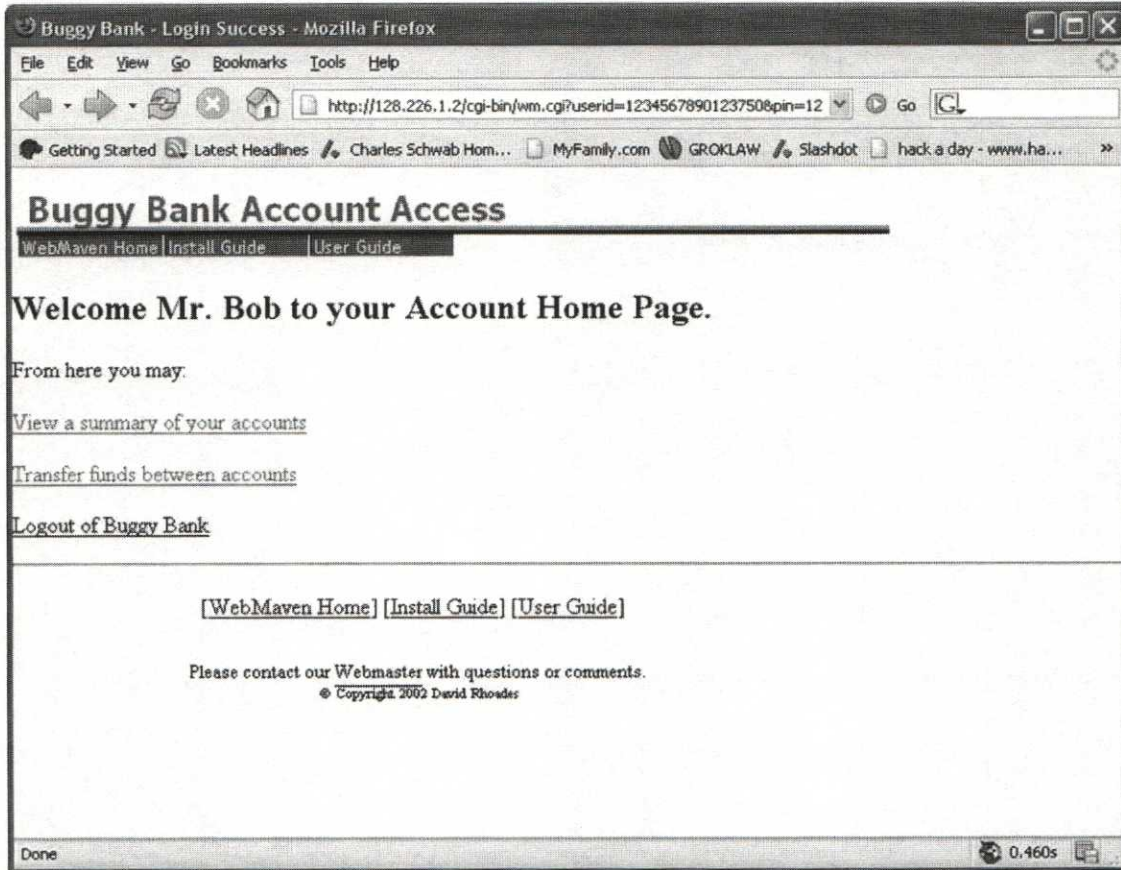


Please notice that there are two account numbers with valid PIN numbers listed on the page. Please highlight and then copy (Control-C) one of these account numbers and

remember the respective PIN number. Once you have done this, please click the "Login" button.

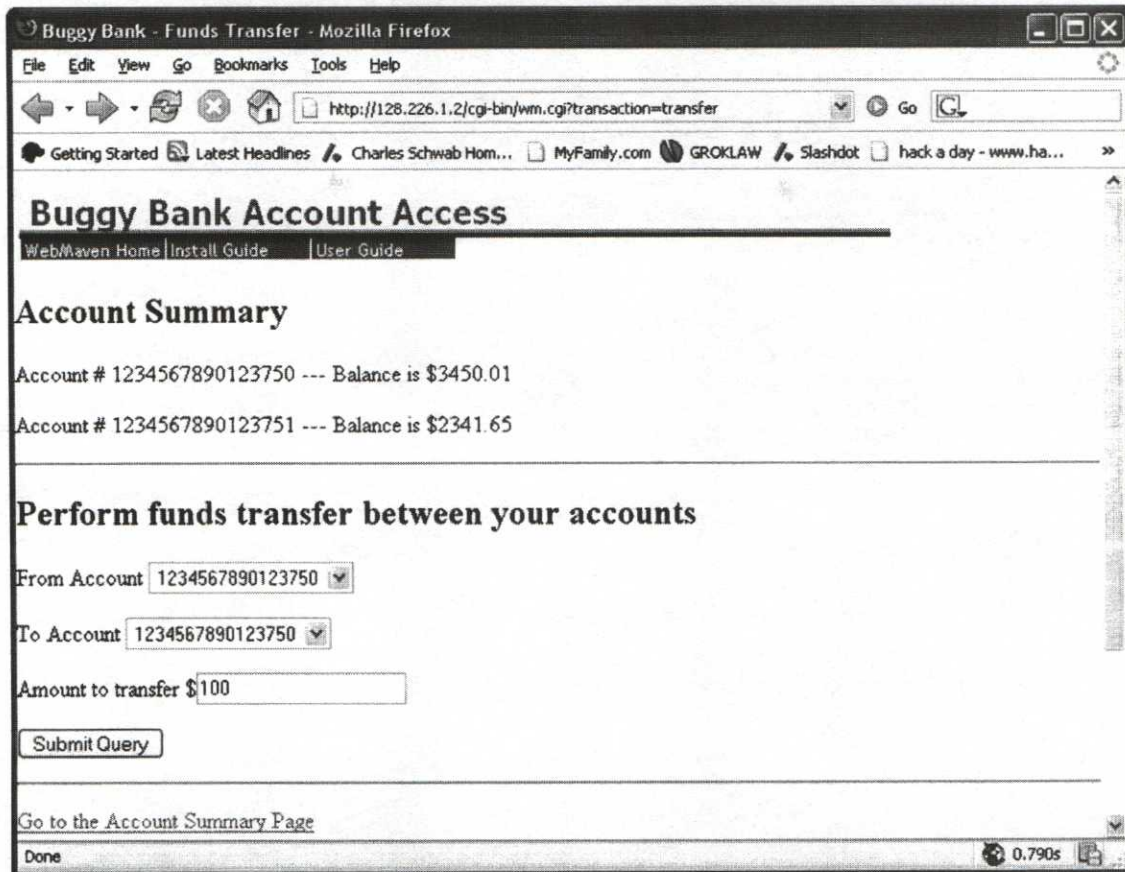


Please paste the account number into the proper box and enter the PIN number that you took note of from the previous page. Once this is done, please click the "Submit Query" button.

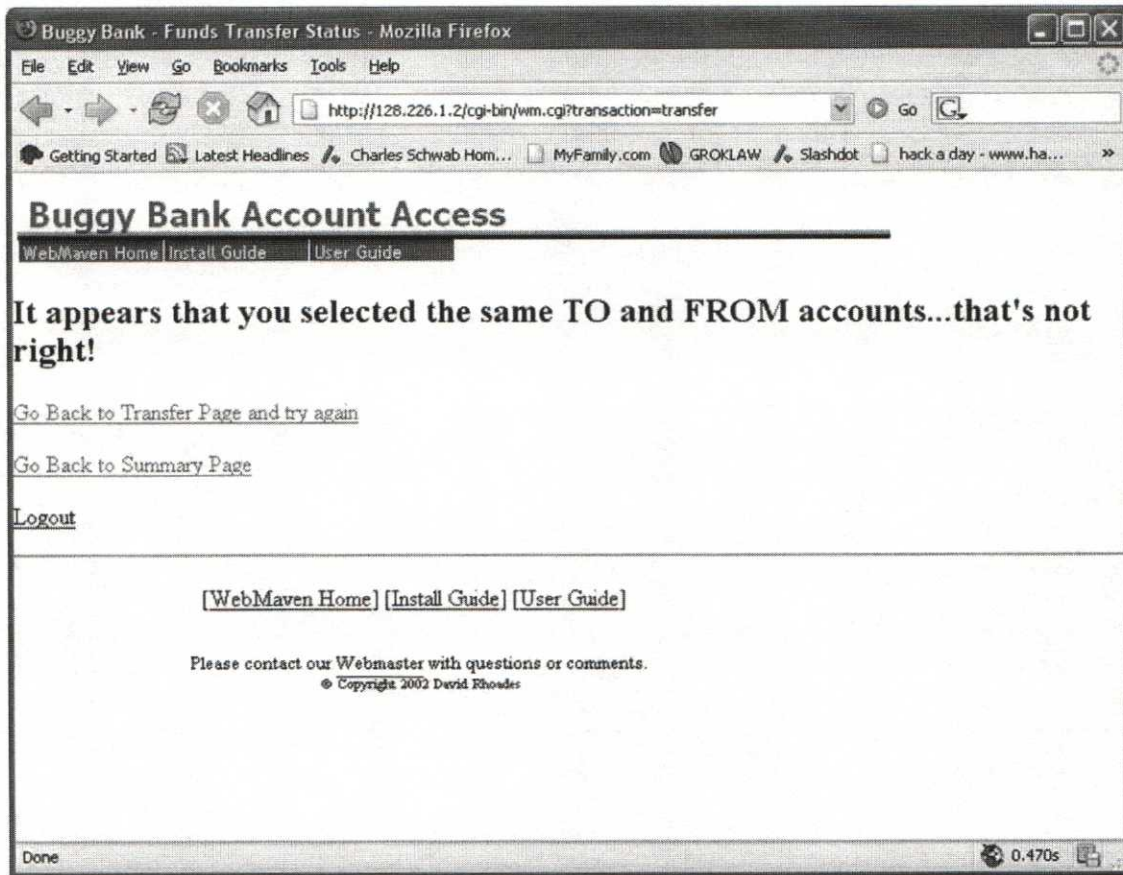


At this point, we have successfully logged in. Although the Burp fuzzer can also be used as a brute-forcer, we will not be using it in this way. Instead, we will be focusing on using it to manipulate input fields.

Please click the “Transfer funds between accounts” option:

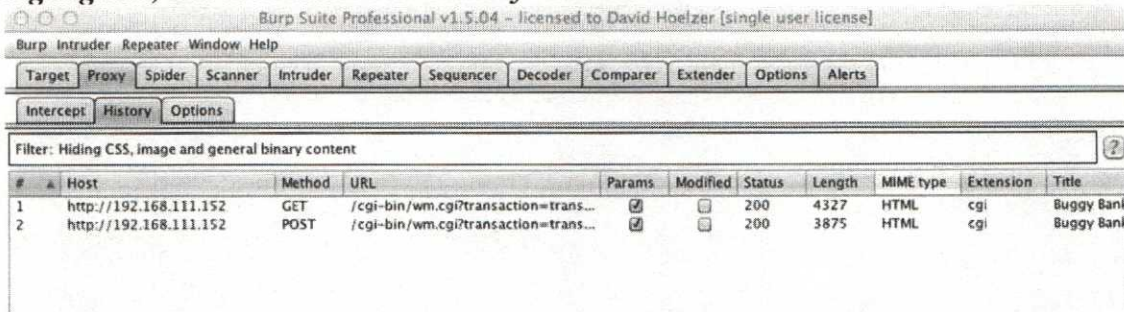


We're still not quite ready. Although it is possible to hand-craft all of the parameters necessary to run the fuzzer, it is much easier to use a previous request as a template. *To accomplish this, please enter any amount of money into the "Amount to transfer" field and click "Submit Query."*



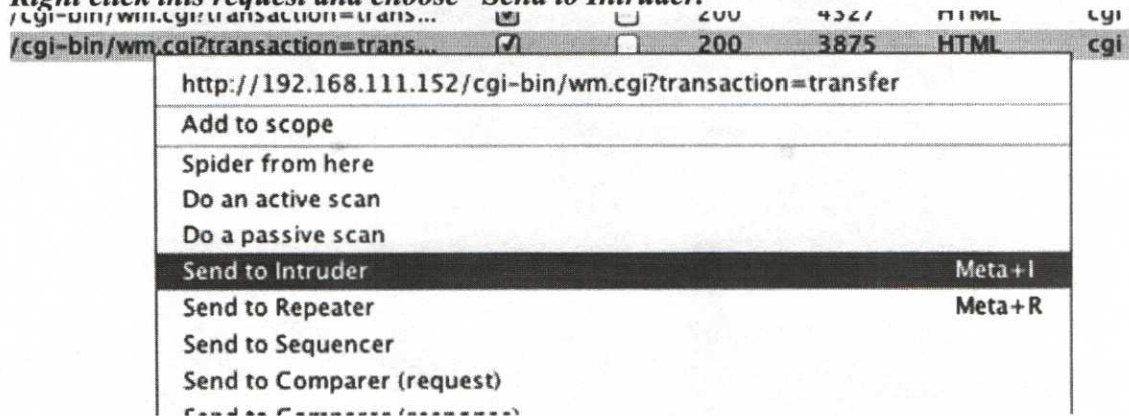
The error that is produced really doesn't matter for us. Remember, we're interested in manipulating the values. Now that we have some, let's take a look at the fuzzer.

Please bring Burp to the foreground. Select the "Proxy" tab if it is not already highlighted, and then select the "History" sub-tab:

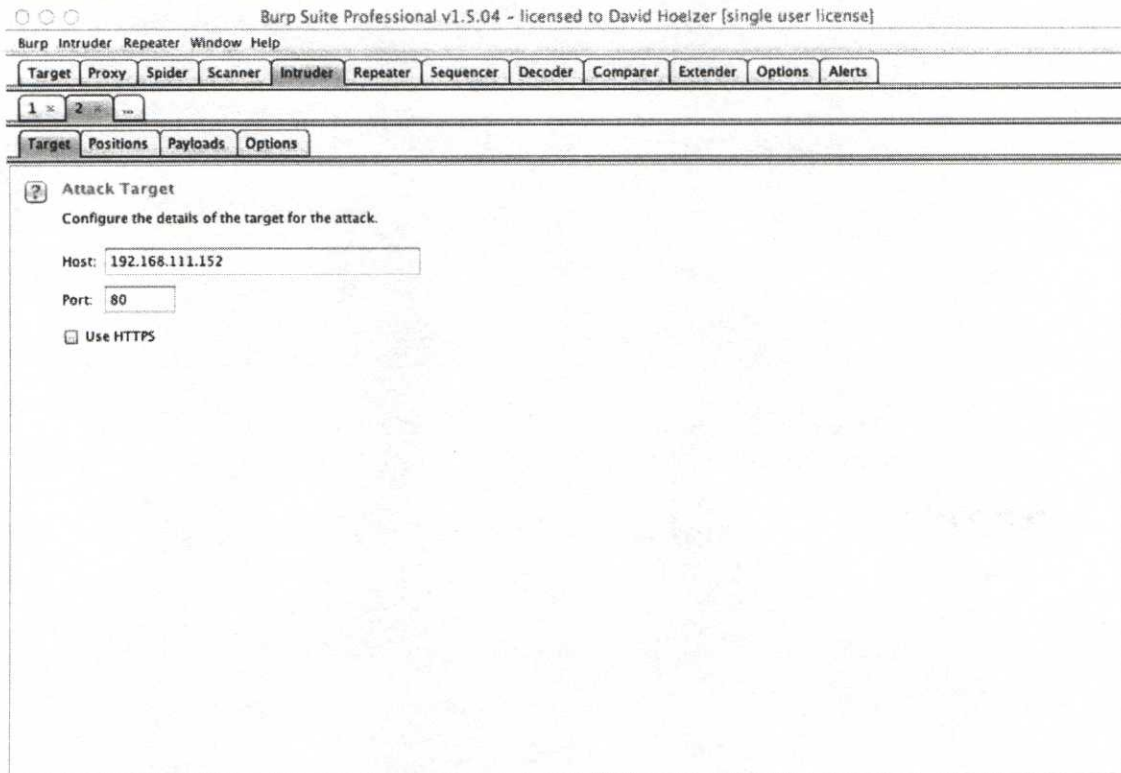


Locate a POST request to the "/cgi-bin/wm.cgi?transaction=transfer..." In the image above, this is the second request.

Right click this request and choose "Send to Intruder."



Please select the "Intruder" tab.



The Intruder is broken into several sections. The first section allows you to configure the server details, which should be correct because we have based this on an actual request.

Please switch to the second tab:

SANS Advanced Systems Audit Workbook

Burp Suite Professional v1.5.04 - licensed to David Hoelzer [single user license]

Target Proxy Spider Scanner Intruder Repeater Sequencer Decoder Comparer Extender Options Alerts

1 2 ...

Target Positions Payloads Options

? Payload Positions

Configure the positions where payloads will be inserted into the base request. The attack type determines the way in which payloads are assigned to payload positions - see help for full details.

Attack type:

```
POST /cgi-bin/wm.cgi?transaction=$transfer$ HTTP/1.1
Host: 192.168.111.152
User-Agent: Mozilla/5.0 (Macintosh; Intel Mac OS X 10.8; rv:15.0) Gecko/20100101
Firefox/15.0.1
Accept: text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8
Accept-Language: en-us,en;q=0.5
Accept-Encoding: gzip, deflate
Proxy-Connection: keep-alive
Referer: http://192.168.111.152/cgi-bin/wm.cgi?transaction=transfer
Cookie: SessionID=$1366295016$
Account=$PCW13MSX88QD3sQfRkKH0m6JCyGca7M6mtaLPn6zLNsmc3l*2FF5FGU10Kg*9Dk3DvV3k*0D$
Content-Type: multipart/form-data;
boundary=-----1057944912119034168738459728
Content-Length: 665
-----1057944912119034168738459728
Content-Disposition: form-data; name="from"
$1234567890123750$
-----1057944912119034168738459728
Content-Disposition: form-data; name="to"
```

0 matches

8 payload positions Length: 1328

As you can see, the Intruder automatically marks everything that it identifies as an input element to the application. At some point in the testing process, you will want to manipulate every element in the form. For now, though, we want to manipulate only the fields that have to do with the actual funds transfer. Why? Well, if we were to manipulate the “SessionID” value, for example, we would no longer be authenticated to the application! This would likely make the testing that we are performing against the transfer function completely invalid.

To make this happen, we must select only the items that we want to test. Let’s start by eliminating all of the current field markers, and then include only the ones that we want.

Please delete the \$ markers from all of the fields. To do this, click the “Clear \$” button to the right. If, when you click the clear button, only one field is cleared, you most likely had inadvertently selected that field. Simply click the “Clear \$” button again to remove all of the markers.

After all of the markers are removed, please highlight the value in the “From” section and click “Add.” Repeat this for the “To” and the “Amount” values as well:

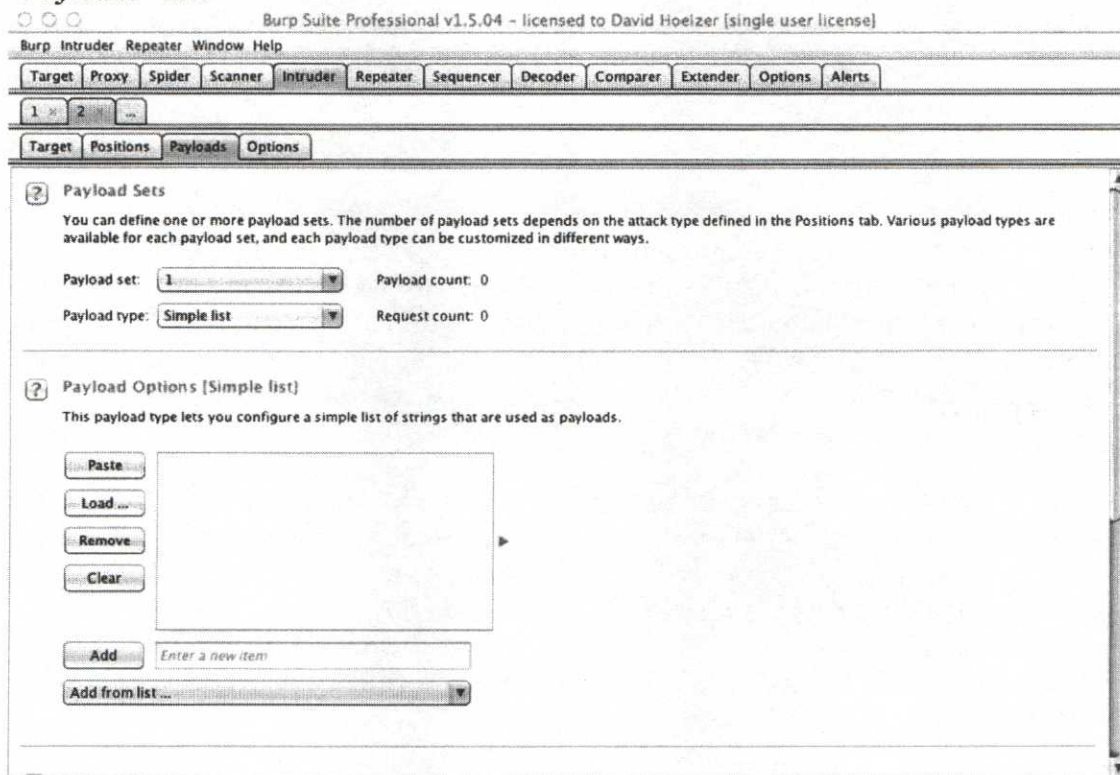
```

POST /cgi-bin/wm.cgi?transaction=transfer HTTP/1.1
Host: 192.168.111.152
User-Agent: Mozilla/5.0 (Macintosh; Intel Mac OS X 10.8; rv:15.0) Gecko/20100101
Firefox/15.0.1
Accept: text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8
Accept-Language: en-us,en;q=0.5
Accept-Encoding: gzip, deflate
Proxy-Connection: keep-alive
Referer: http://192.168.111.152/cgi-bin/wm.cgi?transaction=transfer
Cookie: SessionID=1366285016;
Account=pCqz13mSxE8gD3aQfHeKHOMBJCyGca7M6mtaLPn6zINsSc3l%2FF5FdGUl0Kg%3D%3DvV3i%0D
Content-Type: multipart/form-data;
boundary=-----1057944912119034168738459728
Content-Length: 665

-----1057944912119034168738459728
Content-Disposition: form-data; name="from"

$1234567890123750$
-----1057944912119034168738459728
Content-Disposition: form-data; name="to"
    
```

After selecting the “From,” “To,” and “Amount” fields, please switch to the “Payloads” tab.



The payload manipulation features of Burp are quite detailed. For testing, you are going to stick with a “Simple list.” This will allow you to import a list of strings that should be used to test every field within the form.

Note that you can set multiple payload types and define a separate payload type to use in each form field; you will not go this far in the testing.

Remember that the idea is to jam different types of stuff into every input that the application accepts. To accomplish this, the first thing that we need to do is to create and select a fuzzing source.

Please open the Windows Notepad and enter the following values as samples for fuzzing:

A single quote.

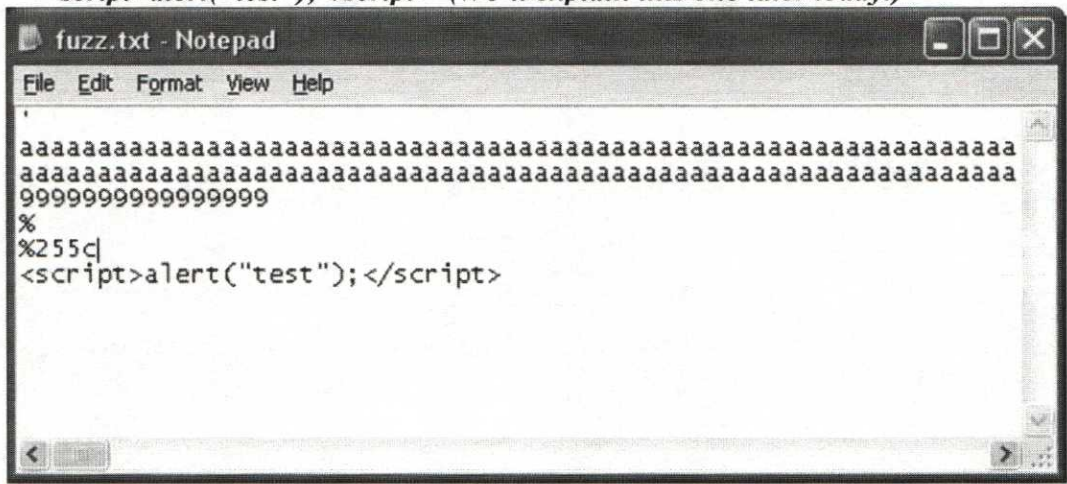
A very long string (several hundred characters... copy and paste works well!)

A very long integer. Many billions work well here.

A lone percent sign.

%255c (to represent an encoded backslash).

<script>alert("test");</script> (We'll explain this one later today.)



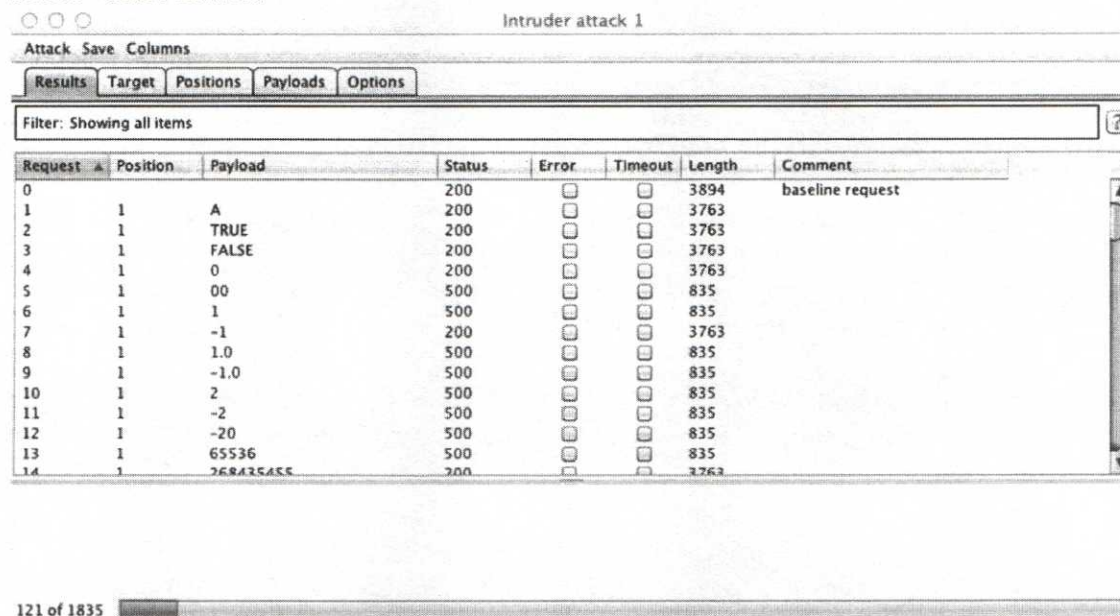
Please save this file on your desktop as "fuzz.txt".

Alternatively, there is a file named "all_attack" in the "Web Application Testing" folder within the "Tools" folder. You might find that your anti-virus tool objects to this file and prevents you from using it, but this file simply contains strings that can be used for testing a web application's responses to bad input.

To set up the list, *click the "Load" button in the "Payload Options – Simple List" section in Burp. Within the window that opens, locate your file of test strings and "Open" it:*



That's it! You're ready to go. *To start the test, pull down the "Intruder" menu and select "Start Attack."*



To analyze the results, start by looking for anomalies from the Baseline Request. Burp starts by sending the request as is and benchmarking it. You can see the status is 200 and the size is 3894 bytes. *(Please note!! You might see a different size!! We received this*

specific size using a value of \$1 to transfer. If you selected \$100, for instance, you might see this baseline vary by a few bytes!)

Note that each of the columns can be sorted. You can begin by sorting by status code. In this test application, there might be many, many queries that will return a 500. In real-world testing, you would expect there to be no 500 errors! This error usually means that you have managed to send some type of input into the web application and that input caused the application to crash! In this case, however, the 500 errors are being caused by the fact that we have a tiny virtual machine with very little memory that is simply unable to handle the volume of requests that are being sent.

In addition, if it is possible to make the application enter an undefined state, which is probably happening when it is heading for a 500 response code, it is entirely possible that this indicates that there might be something exploitable!

The next column to examine is the Length column. For instance, if you found that the correct transaction results in 3800 bytes or so and that an error results in about 3700 bytes or so, sorting by size allows you to find *other* sizes! These all indicate that the application processed the data in an unexpected way and returned what might be very interesting data!

Of course, as you examine these and find something interesting, double-clicking any line in the Intruder window will open the transaction for you, allowing you to determine precisely which input values caused the problem!

On your own, see whether you can successfully configure the fuzzer to run against the NLog webpage located on the web server VM. You should have recorded the IP address for this earlier today.

Knowing how to use the fuzzer can be a real time saver when performing an audit. Prior to the creation of fuzzers of this sort, testing required hours of manual time spent testing combinations that would most likely produce interesting results. With the fuzzer, though, you can set things up and let them run without intervention, and then review the results at your leisure.

Exercise 3: Brute-Force Authentication/Credential Exposure

Purpose: Demonstrate the vulnerability inherent in username/password authentication. Demonstrate the ability to distinguish valid accounts when controls are not consistent.

Username and passwords are typically used as primary authentication mechanisms to allow access. In other words, they are most usually the controls that are in place to meet access control objectives with regard to information. Although this is not necessarily a bad thing (far less expensive than biometric authentication, for instance), it is not the most secure mechanism unless we introduce additional controls. For instance, if we were to secure access to portions of a webapp or website based on username and password authentication, we will want additional controls, like account lockouts, in place to ensure that the data is not compromised.

Configuring Burp

To get started, you must first configure the browser to run all of its queries through Burp. *Please verify that your browser is currently configured to proxy traffic through Burp.*

With the proxy configured, all that we have to do is attempt to log in and then review the activity with Burp. *Please go to the bank login page and attempt to login with invalid data.*

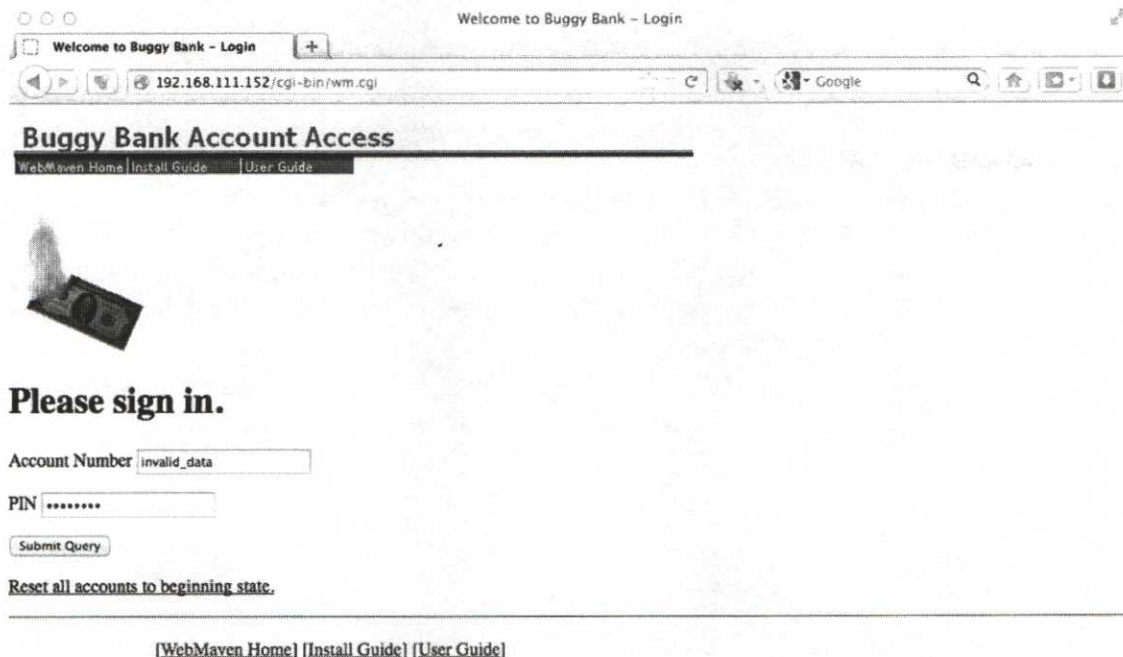


Figure 3 - Invalid Data Entered

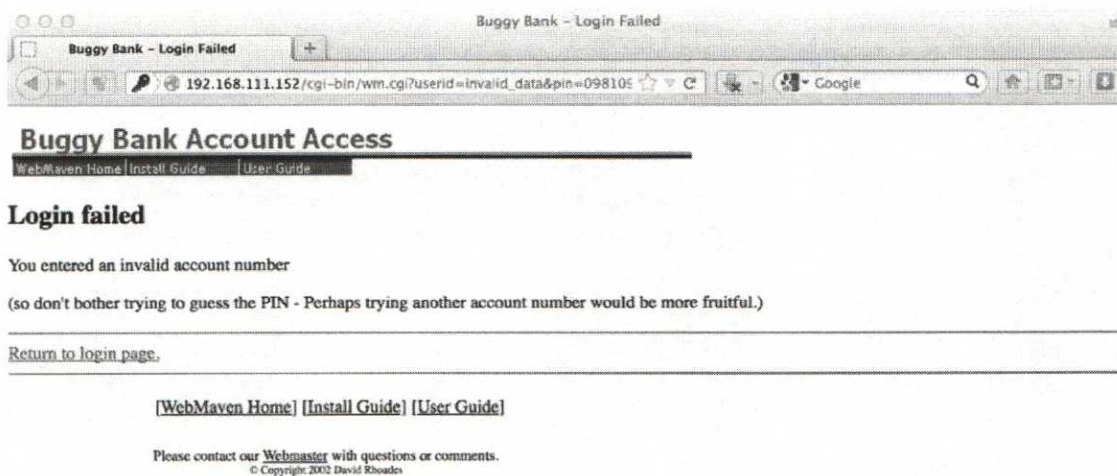


Figure 4 - Failed login with invalid data

Now that we failed to log in, we can use the information in the history logs to create a brute-force attack. *Within Burp, switch to the “Proxy” tab, and then select the “History” tab.*

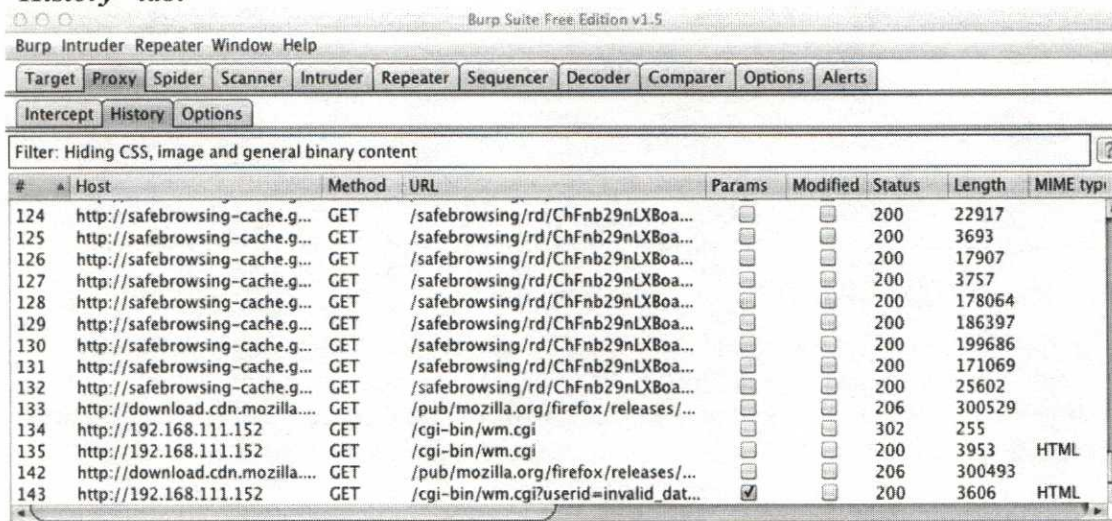


Figure 5 - Transaction history from Burp

Here we can see all of the past transactions. If you have been following along, it is quite likely that the last transaction in your history, like the one above, is the login attempt. **Select the last entry and verify that it is, in fact, the failed login attempt.**

Burp Suite Free Edition v1.5

Burp Intruder Repeater Window Help

Target Proxy Spider Scanner Intruder Repeater Sequencer Decoder Comparer Options Alerts

Intercept History Options

Filter: Hiding CSS, image and general binary content

#	Host	Method	URL	Params	Modified	Status	Length	MIME type
124	http://safebrowsing-cache.g...	GET	/safebrowsing/rd/ChFnb29nLXBoa...	<input type="checkbox"/>	<input type="checkbox"/>	200	22917	
125	http://safebrowsing-cache.g...	GET	/safebrowsing/rd/ChFnb29nLXBoa...	<input type="checkbox"/>	<input type="checkbox"/>	200	3693	
126	http://safebrowsing-cache.g...	GET	/safebrowsing/rd/ChFnb29nLXBoa...	<input type="checkbox"/>	<input type="checkbox"/>	200	17907	
127	http://safebrowsing-cache.g...	GET	/safebrowsing/rd/ChFnb29nLXBoa...	<input type="checkbox"/>	<input type="checkbox"/>	200	3757	
128	http://safebrowsing-cache.g...	GET	/safebrowsing/rd/ChFnb29nLXBoa...	<input type="checkbox"/>	<input type="checkbox"/>	200	178064	
129	http://safebrowsing-cache.g...	GET	/safebrowsing/rd/ChFnb29nLXBoa...	<input type="checkbox"/>	<input type="checkbox"/>	200	186397	
130	http://safebrowsing-cache.g...	GET	/safebrowsing/rd/ChFnb29nLXBoa...	<input type="checkbox"/>	<input type="checkbox"/>	200	199686	
131	http://safebrowsing-cache.g...	GET	/safebrowsing/rd/ChFnb29nLXBoa...	<input type="checkbox"/>	<input type="checkbox"/>	200	171069	
132	http://safebrowsing-cache.g...	GET	/safebrowsing/rd/ChFnb29nLXBoa...	<input type="checkbox"/>	<input type="checkbox"/>	200	25602	
133	http://download.cdn.mozilla...	GET	/pub/mozilla.org/firefox/releases/...	<input type="checkbox"/>	<input type="checkbox"/>	206	300529	
134	http://192.168.111.152	GET	/cgi-bin/wm.cgi	<input type="checkbox"/>	<input type="checkbox"/>	302	255	
135	http://192.168.111.152	GET	/cgi-bin/wm.cgi	<input type="checkbox"/>	<input type="checkbox"/>	200	3953	HTML
142	http://download.cdn.mozilla...	GET	/pub/mozilla.org/firefox/releases/...	<input type="checkbox"/>	<input type="checkbox"/>	206	300493	
143	http://192.168.111.152	GET	/cgi-bin/wm.cgi?userid=invalid_dat...	<input checked="" type="checkbox"/>	<input type="checkbox"/>	200	3606	HTML

Request Response

Raw Headers Hex HTML Render

Buggy Bank Account Access

WebWaver Home | Install Guide | User Guide

Login failed

You entered an invalid account number

(so don't bother trying to guess the PIN - Perhaps trying another account number would be more fruitful.)

[Return to login page](#)

Figure 6 - Selecting the failed attempt

With the failure selected, you can see in Figure 6 that we have additionally selected the Response section and asked it to “Render” the returned value. This allows us to see the actual failure page in the bottom pane. **After you have verified that you have found the login failure, right-click the request in the top pane and select “Send to Intruder.”**

The Intruder will allow us to configure Burp to run any of a number of attacks against the login form. To get started, please select the “Intruder” tab, which should have turned orange when you sent the request to the Intruder.

SANS Advanced Systems Audit Workbook

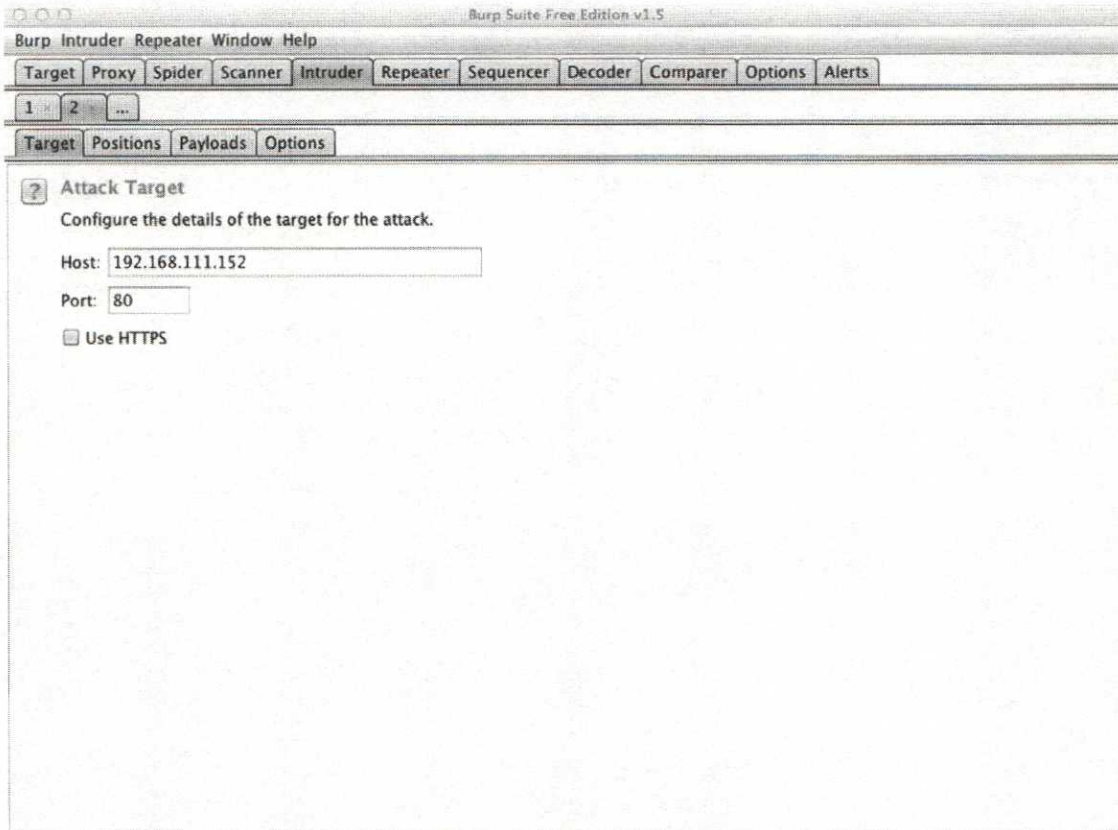


Figure 7 - Selecting the Intruder tab

To use the Intruder, we have just a few things that we need to configure. First, you should see that the “Target” tab is already correctly configured based on the request that we had previously selected.

Select the “Positions” tab.

SANS Advanced Systems Audit Workbook

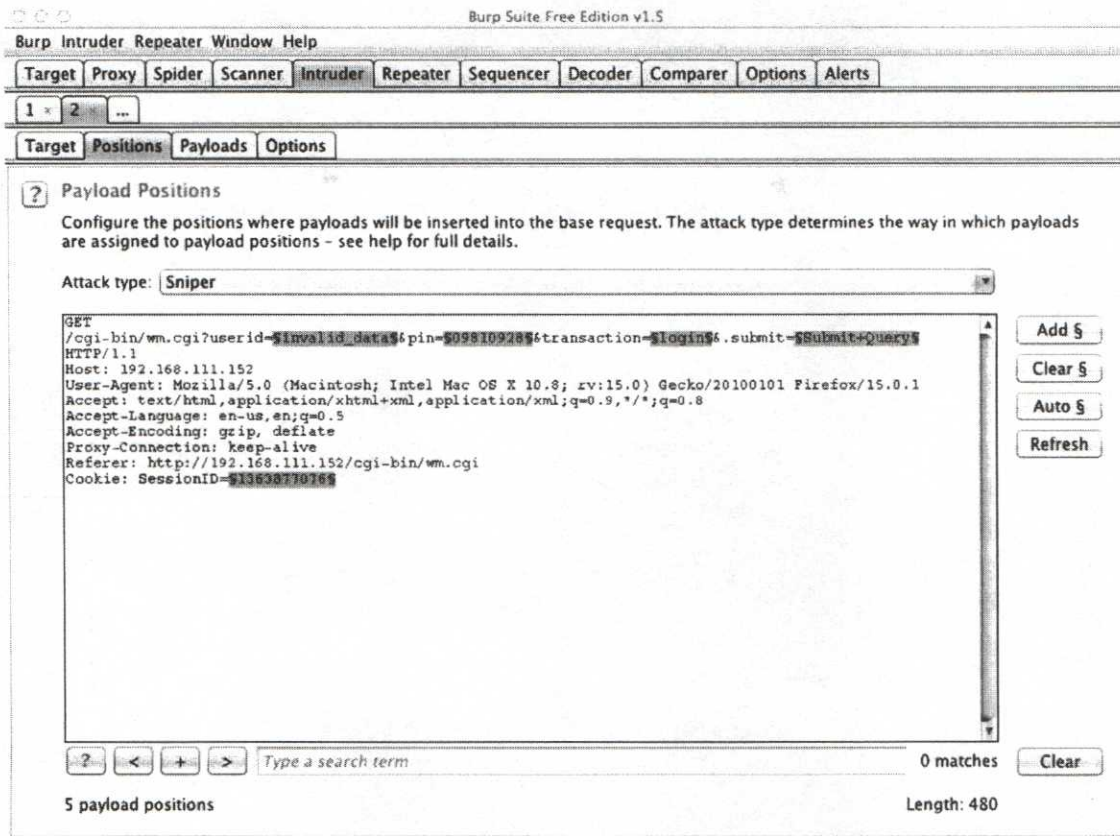


Figure 8 - Intruder Positions

With positions selected, we can now tell Burp which fields we'd like to automatically replace. You should notice that, by default, it will select every input field that's available. *For our purposes, because we simply want to attempt to brute force an account, let's start by clicking the "Clear" button:*




Figure 9 - The Clear button

Once you have selected "Clear," all of the highlighted portions will disappear. *Next, locate the "pin" field within the submitted URL. Highlight the value that appears there and click the "Add" button.*

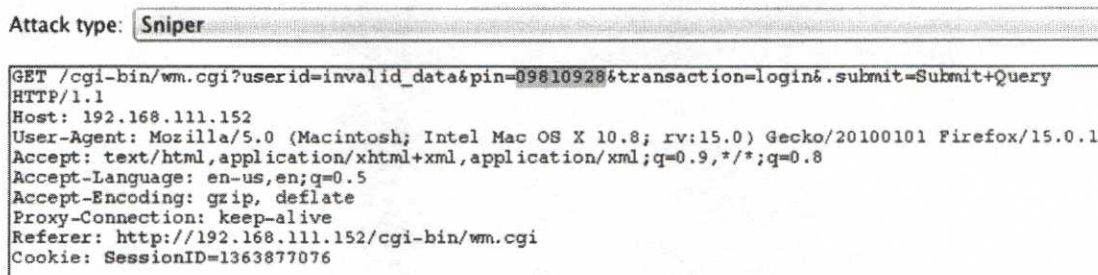


Figure 10 - Selecting the pin value

```

id_data&pin=$09810928&&transaction=log:
;osh; Intel Mac OS X 10.8; rv:15.0) Gecko
(html+xml,application/xml;q=0.9,*/*;q=0.8
;

```

Figure 11 - The pin value after clicking "Add"

Now that we have Burp set up to inject data into the pin parameter, we need to tell it what kind of data to insert. *Please click the "Payloads" tab.*

The Payloads section allows you to customize how the injection attack will run. We know based on information provided that this application uses a four-digit PIN # for authentication. This allows us to focus our test.

Pull down the "Payload Type" menu and select "Numbers."

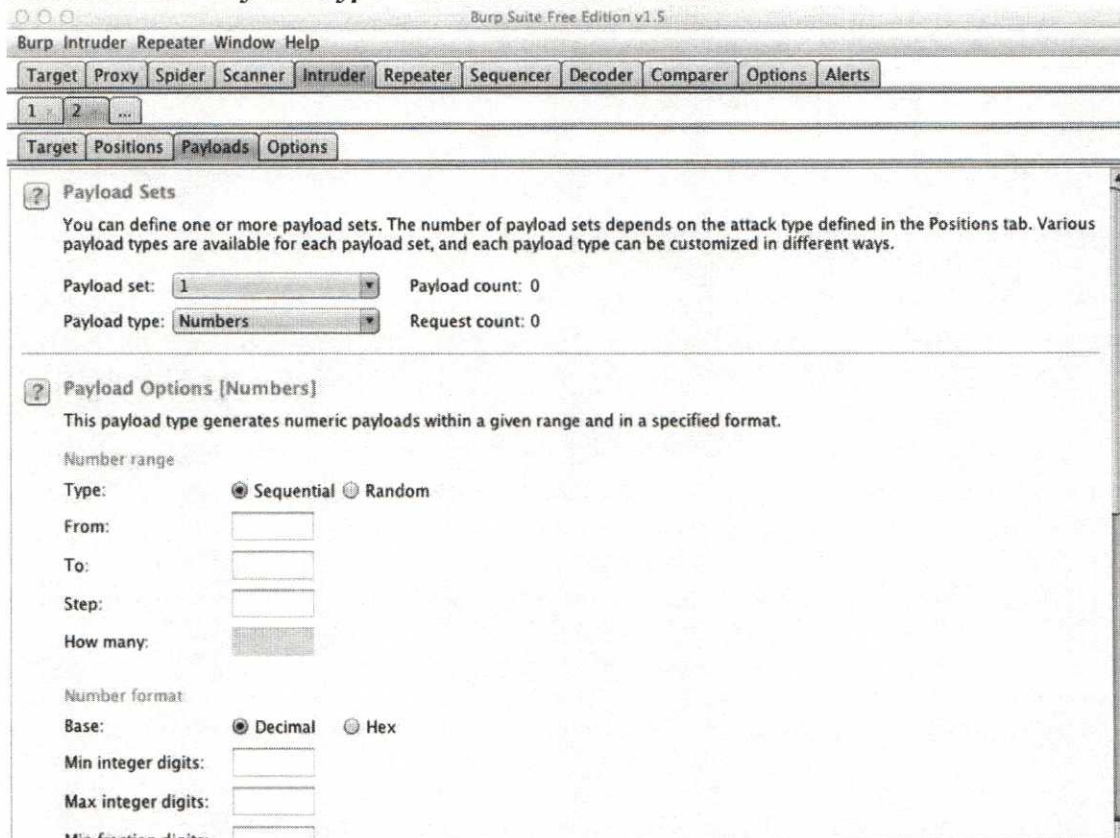



Figure 12 - Payload type set to numbers

Next, set the system to use sequential numbers from 0000 to 9999 stepping by 1. In addition, set the minimum and maximum number of integer digits to 4 and the minimum and maximum fraction digits to zero:

 Payload Options [Numbers]

This payload type generates numeric payloads within a given range and in a specified format.

Number range

Type: Sequential Random

From:

To:

Step:

How many:

Number format

Base: Decimal Hex

Min integer digits:

Max integer digits:

Min fraction digits:

Max fraction digits:

Examples

0001

4321

Figure 13 - Configuring payload options

There are additional options that can be configured, but for basic testing this is really all that is required. In fact, this is precisely how fuzzing tests are configured in Burp except that rather than selecting a list of numbers, we would likely import data from a file.

Pull down the “Intruder” menu and select “Start attack.” When you do so, a warning will appear:




Figure 14 - Intruder Warning

Because we are using the free version of Burp in this class, there are a few limitations. One of them is that it will not run an intruder test full speed. For our purposes, this isn't a problem. We just need to see how it functions and also see whether we can distinguish a valid username from an invalid username.

Click **OK** to start the test.

The screenshot shows a window titled "Intruder: attack 1". At the top, there are buttons for "Attack", "Save", and "Columns". Below that are tabs for "Results", "Target", "Positions", "Payloads", and "Options". A filter bar indicates "Showing all items". The main area contains a table with the following data:

Request	Payload	Status	Error	Timeout	Length	Comment
0		200	<input type="checkbox"/>	<input type="checkbox"/>	3625	baseline request
1	0000	200	<input type="checkbox"/>	<input type="checkbox"/>	3625	
2	0001	200	<input type="checkbox"/>	<input type="checkbox"/>	3625	
3	0002	200	<input type="checkbox"/>	<input type="checkbox"/>	3625	
4	0003	200	<input type="checkbox"/>	<input type="checkbox"/>	3625	
5	0004	200	<input type="checkbox"/>	<input type="checkbox"/>	3625	
6	0005	200	<input type="checkbox"/>	<input type="checkbox"/>	3625	
7	0006	200	<input type="checkbox"/>	<input type="checkbox"/>	3625	
8	0007	200	<input type="checkbox"/>	<input type="checkbox"/>	3625	
9	0008	200	<input type="checkbox"/>	<input type="checkbox"/>	3625	
10	0009	200	<input type="checkbox"/>	<input type="checkbox"/>	3625	

At the bottom of the window, there is a progress bar showing "9 of 10000".

9 of 10000

Figure 15 - Intruder running

You should see an additional window pop open in which the intruder will now run. ***There is no need to allow this to run for a long period of time!!!*** After 10 or 20 seconds, stop the test.

Please review any one of the requests and responses. Notice that they are all exactly the same size. In fact, if you were to review the content, you would find that every response is completely identical.

At this point, we can see how this can be used for brute-force password testing in addition to fuzzing. Can we use it to discern valid from invalid credentials and test password lockouts? Yes! Here's how.

In your web browser, go back to the main page for the bank.

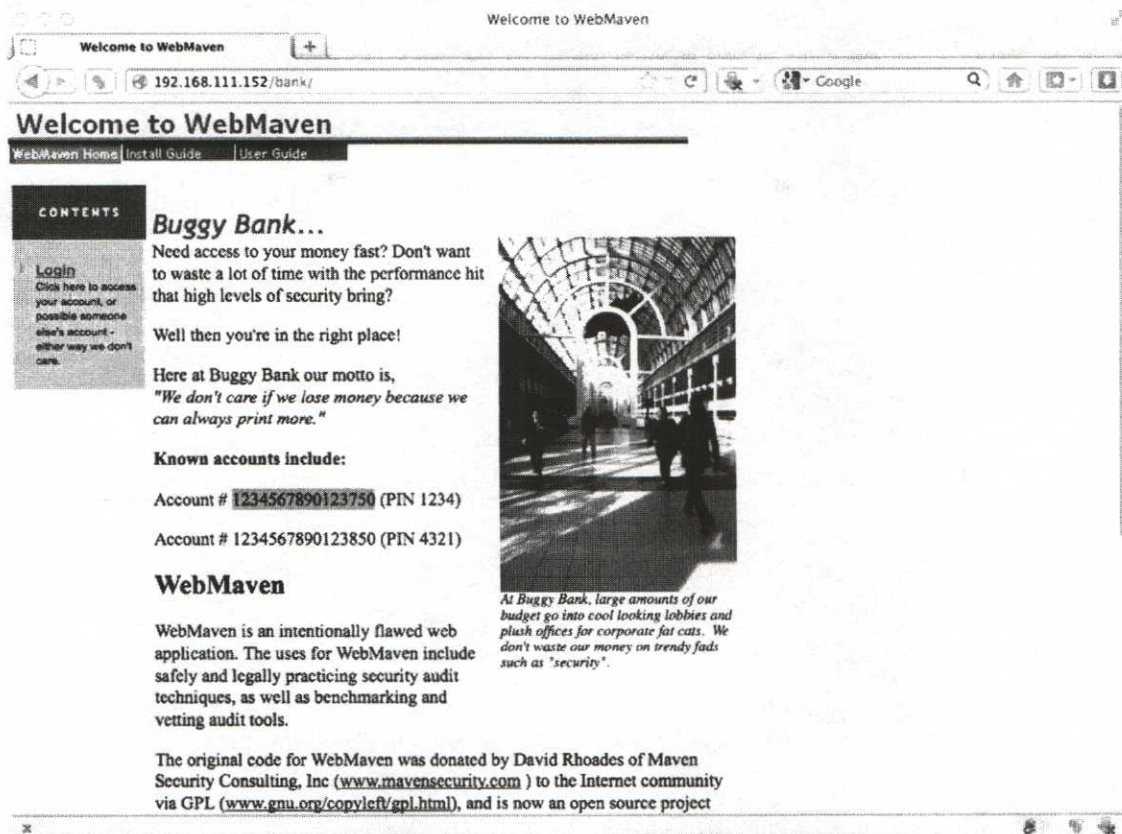


Figure 16 - Bank main page

On the main page, you will see that there is account information for two different accounts. **Select the account number for either one of them and copy it.** It is very important to make sure that you have not included an extra space in what you have copied.

With this data in your clipboard, switch back to Burp. **Close the Intruder window that is currently open. You will be warned that this will exit the current attack. Click "OK."**

You should now be back at the Burp "Payload" window. **Click the "Positions" tab. Within the positions tab, locate the "userid" parameter and change whatever value is assigned to be the account number that you copied in the last step.**

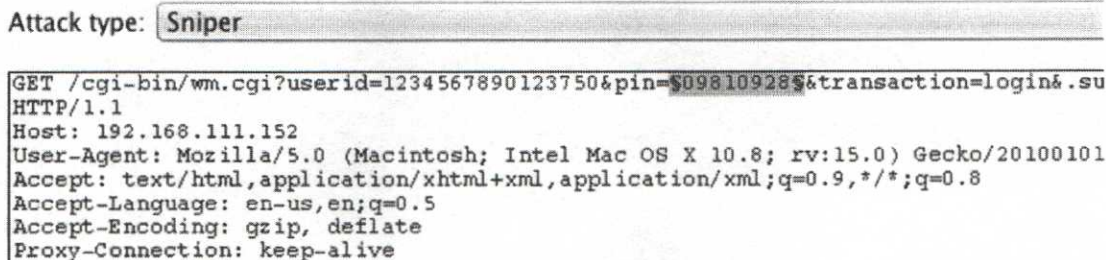


Figure 17 - Userid replaced with a valid account number

Once again, start the attack by pulling down the Intruder menu and selecting "Start attack."

After this runs for just a few seconds, you should see something quite different from the last test. *After the third attempt or so, you should see something change in the output. It will likely look something like this:*

Status	Error	Timeout	Length	Comment
200	<input type="checkbox"/>	<input type="checkbox"/>	3664	baseline request
200	<input type="checkbox"/>	<input type="checkbox"/>	3664	
200	<input type="checkbox"/>	<input type="checkbox"/>	3664	
200	<input type="checkbox"/>	<input type="checkbox"/>	3584	
200	<input type="checkbox"/>	<input type="checkbox"/>	3584	
200	<input type="checkbox"/>	<input type="checkbox"/>	3584	
200	<input type="checkbox"/>	<input type="checkbox"/>	3584	
200	<input type="checkbox"/>	<input type="checkbox"/>	3584	

Figure 18 - A change occurs!

1. *What has changed?*
 2. *Are there differences in the actual response data?*
 3. *Why would this be a finding when compared to the responses with invalid data?*
1. The size of the response has changed.
 2. Yes. The response had been reporting a bad PIN. It then begins to report that the account has been locked out.
 3. Two reasons. First, invalid data reports "Bad account name" whereas a valid account number with an invalid PIN reports "Bad pin." This allows an attacker to easily discover valid accounts. Secondly, only valid accounts lock out, and when they do, they report it, giving an attacker another mechanism for discovering valid accounts.

Exercise 4: Session Analysis

Purpose: To demonstrate the configuration and use the Session Analysis feature of WebScarab and Burp Suite.

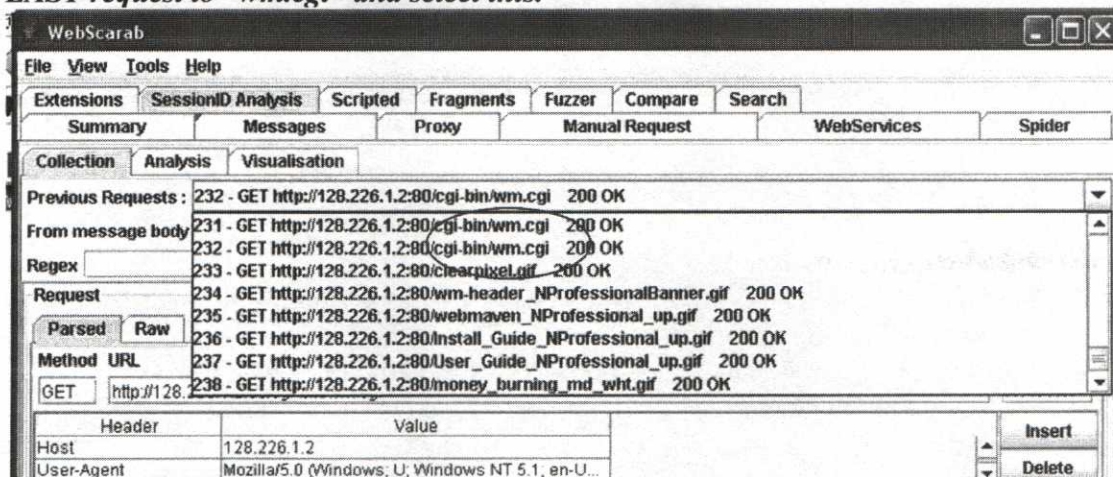
WebScarab

This exercise goes by pretty quickly, so if you don't pay attention you might miss it! Essentially, all that we want to do is collect lots of new session IDs and get WebScarab to graph them out for us so that we can see whether they appear to be generated in a random way.

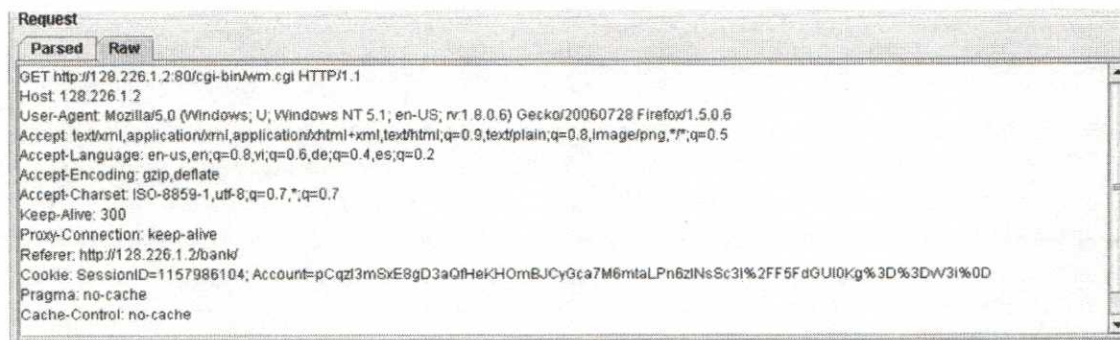
To begin, please return to the BuggyBank homepage at <http://xxx.xxx.xxx.xxx/bank/>. Once you are there, please click the "Login" link. Please force a refresh of this page by holding down the Shift key and pressing the Refresh button.

Please switch to your WebScarab window and select the "Session ID Analysis" tab.

We need to select the last request that was sent to the "wm.cgi" application to get started. *Pull down the "Previous Requests" drop-down menu and scroll to the bottom. Find the LAST request to "wm.cgi" and select this.*



Please make the window large enough that you can clearly see all of the sections of the window. Also, please click the "Raw" tab under the "Request" section. When you do this, you will notice that within the request that was sent, there is a "Cookie" field:

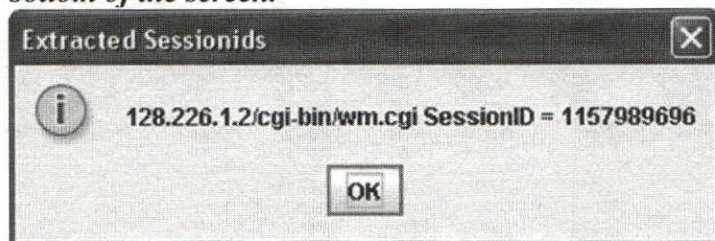


At this point in our interaction with the web application, a session ID has already been generated and set on the client. What you will do is remove this cookie so that the web application will automatically generate a new one!

Please select and delete the “Cookie: SessionID=....” line from the request window.

By deleting this line and resending this request, you will force the application to generate a new session ID. ***PLEASE NOTE:*** This will not always be the case! This just happens to be how this application functions. To perform session analysis, you will need to review the log of transactions and determine at what point that the session ID is created and sent to the client. That is the request that you would use to perform this same task.

After the “Cookie” line has been completely deleted, click the “Test” button at the bottom of the screen.



If the box above appears, you are ready to go! If it does not, please call over a proctor or the instructor for assistance, or just restart the exercise from the beginning!


Now that we have successfully retrieved a session ID, we are ready to graph out a series of them to test for randomness. ***Please set the “Samples” value at the bottom of the page to 250. This will determine how many different session IDs to retrieve. More is better, of course, but 250 should be plenty for our purposes.***

After you have set this value, please click the “Fetch” button.

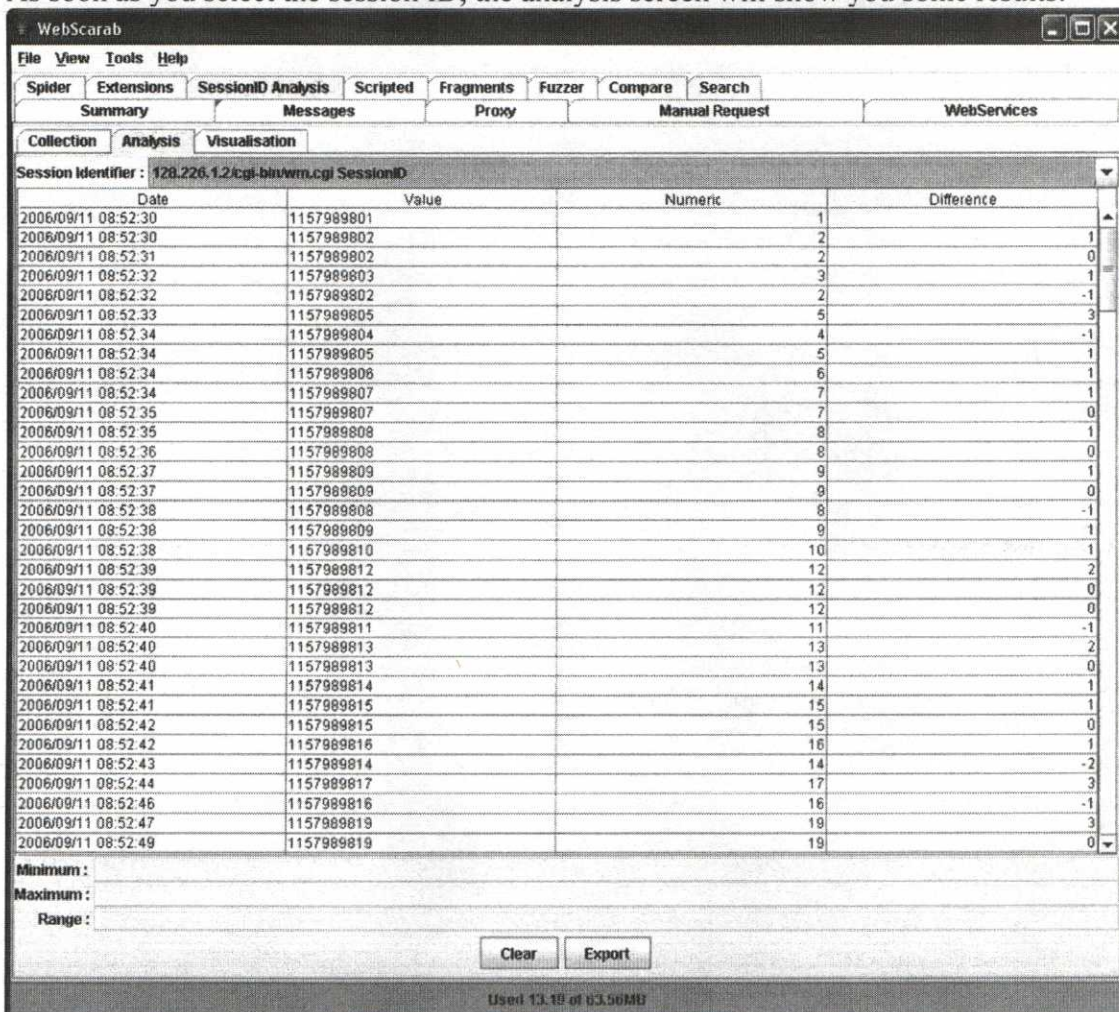
You might be dismayed that clicking the Fetch button has no immediate or apparent result. Don’t despair! The results are being generated even now.

Please notice near the top of the page that there are three tabs available for Session Analysis: Collection, Analysis and Visualization. Please select the "Analysis" tab.

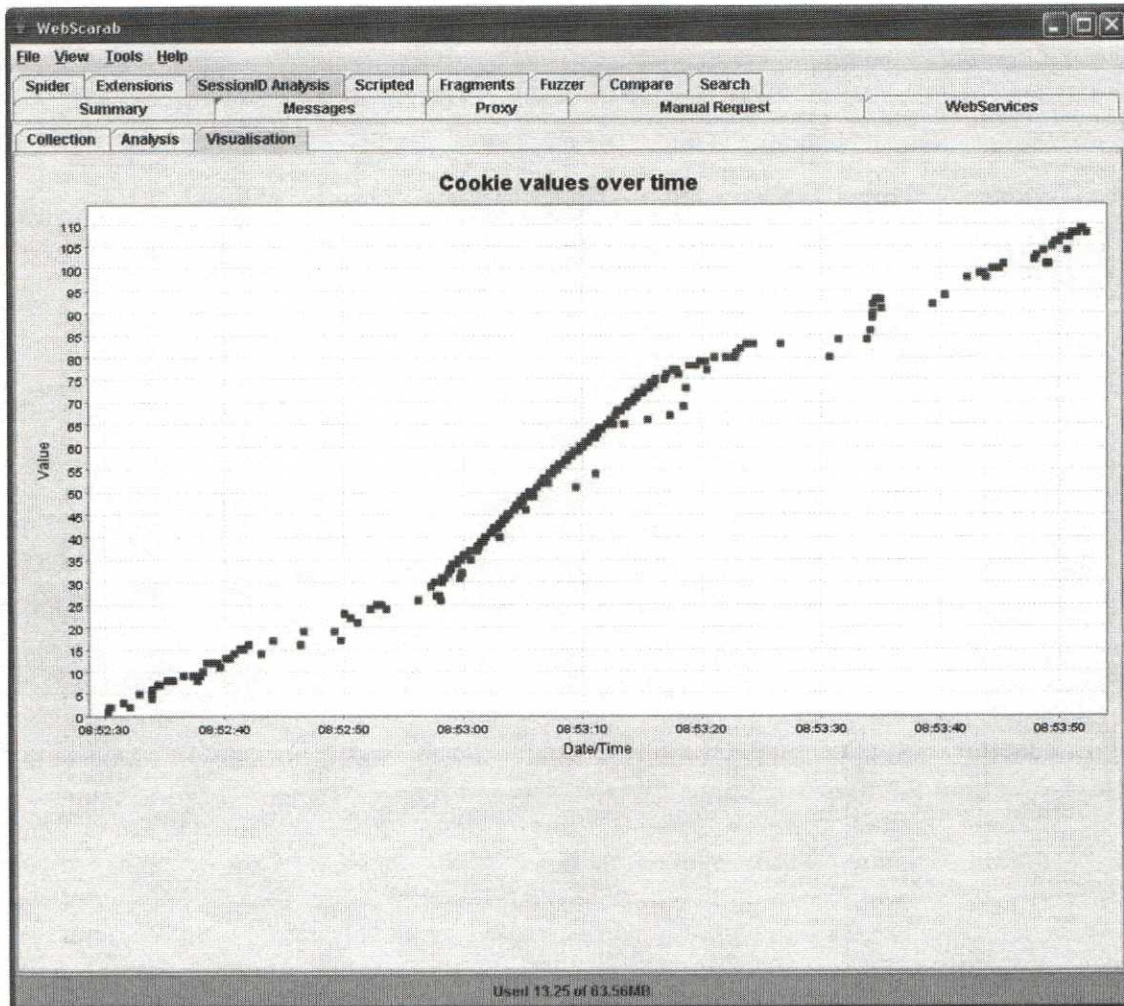
Within the analysis tab, we can review the analysis results. Nothing should be visible when you first select this tab, but if you look at the top of this section, you will notice that there is a pull-down menu labeled "Session Identifier." Pull this down and select the session ID that we were working with. This should be the only option available.



As soon as you select the session ID, the analysis screen will show you some results:



Although these results are interesting and can be used for analysis, sometimes you have really complicated looking session IDs. For this reason, the visualization function is especially useful. Please click the "Visualization" tab.



Looking at the results of visualizing changing session IDs over time, would you say that the session IDs being used are “Strong”?

Why do you say this?

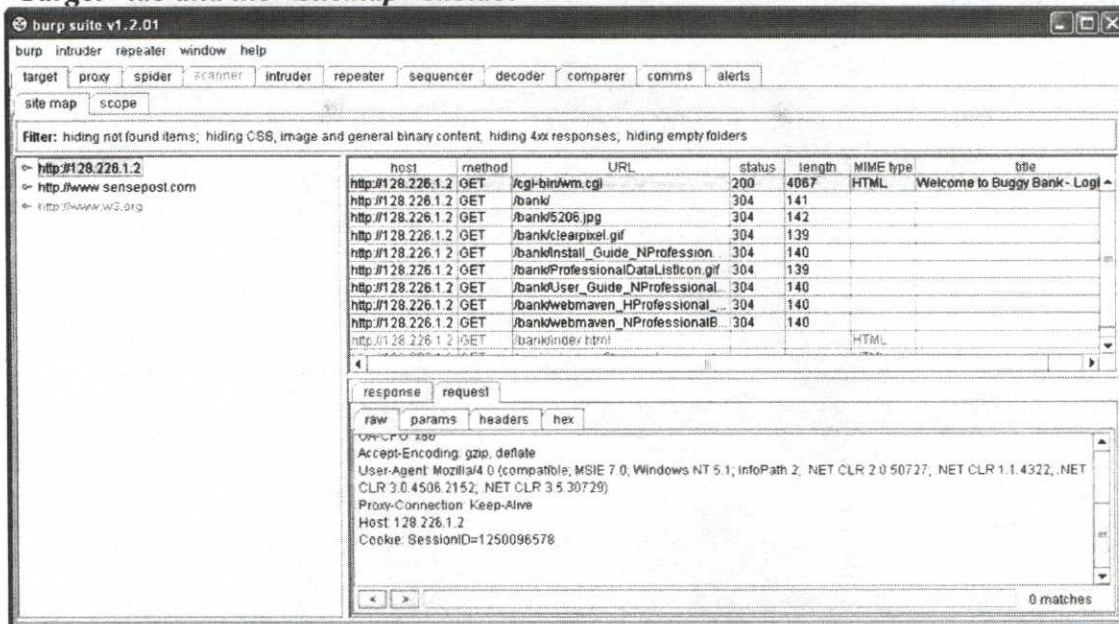
What would you expect the graph of a “Strong” session ID to look like?

Burp Suite

We will now repeat this exercise using Burp Suite.

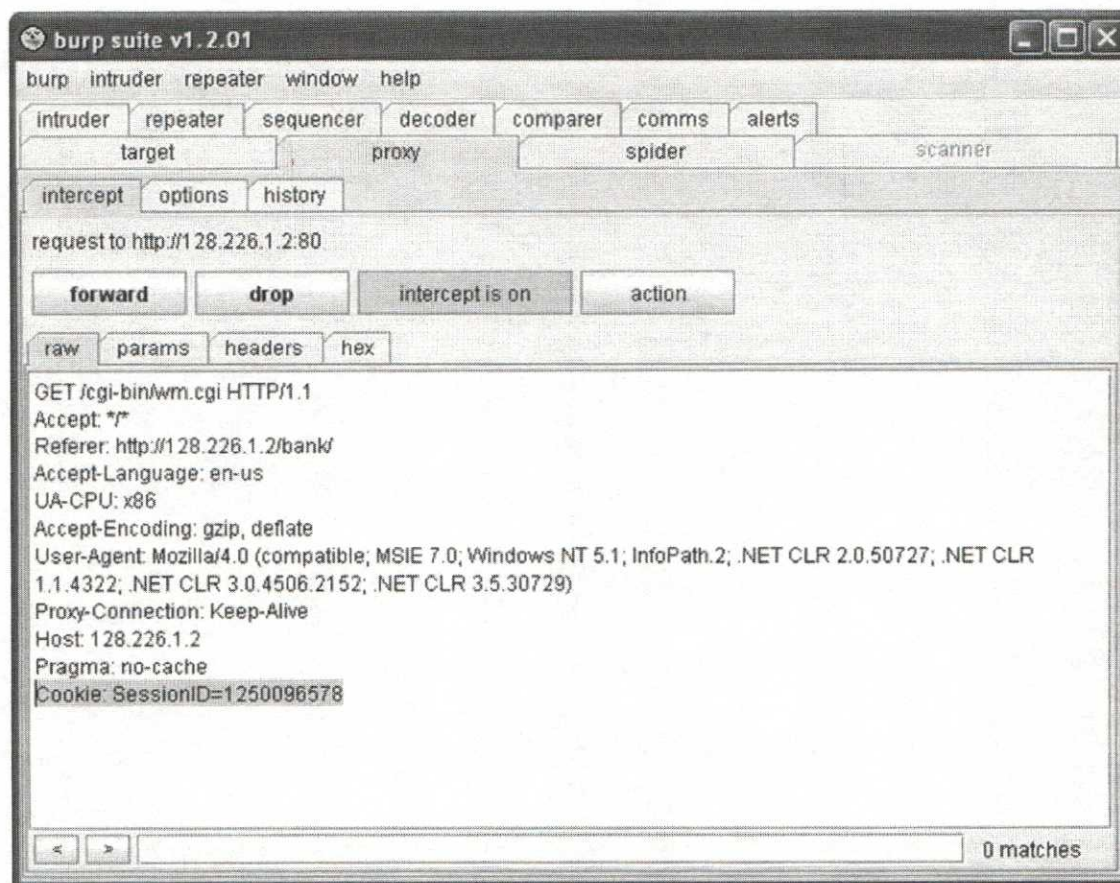
To begin, make sure that your browser is configured to use Burp as a proxy. After that is done, go to the logon page for BuggyBank.

After visiting the logon page, please switch to the Burp Suite window. Choose the “Target” tab and the “Sitemap” subtab:



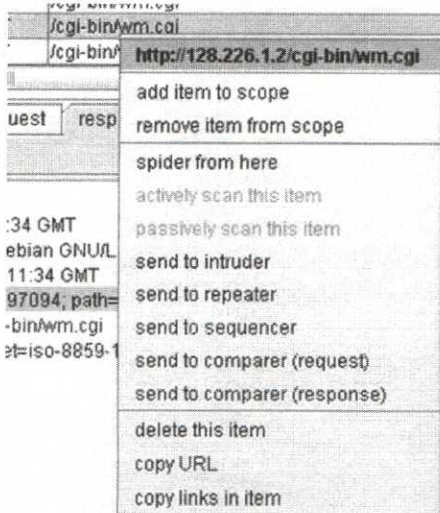
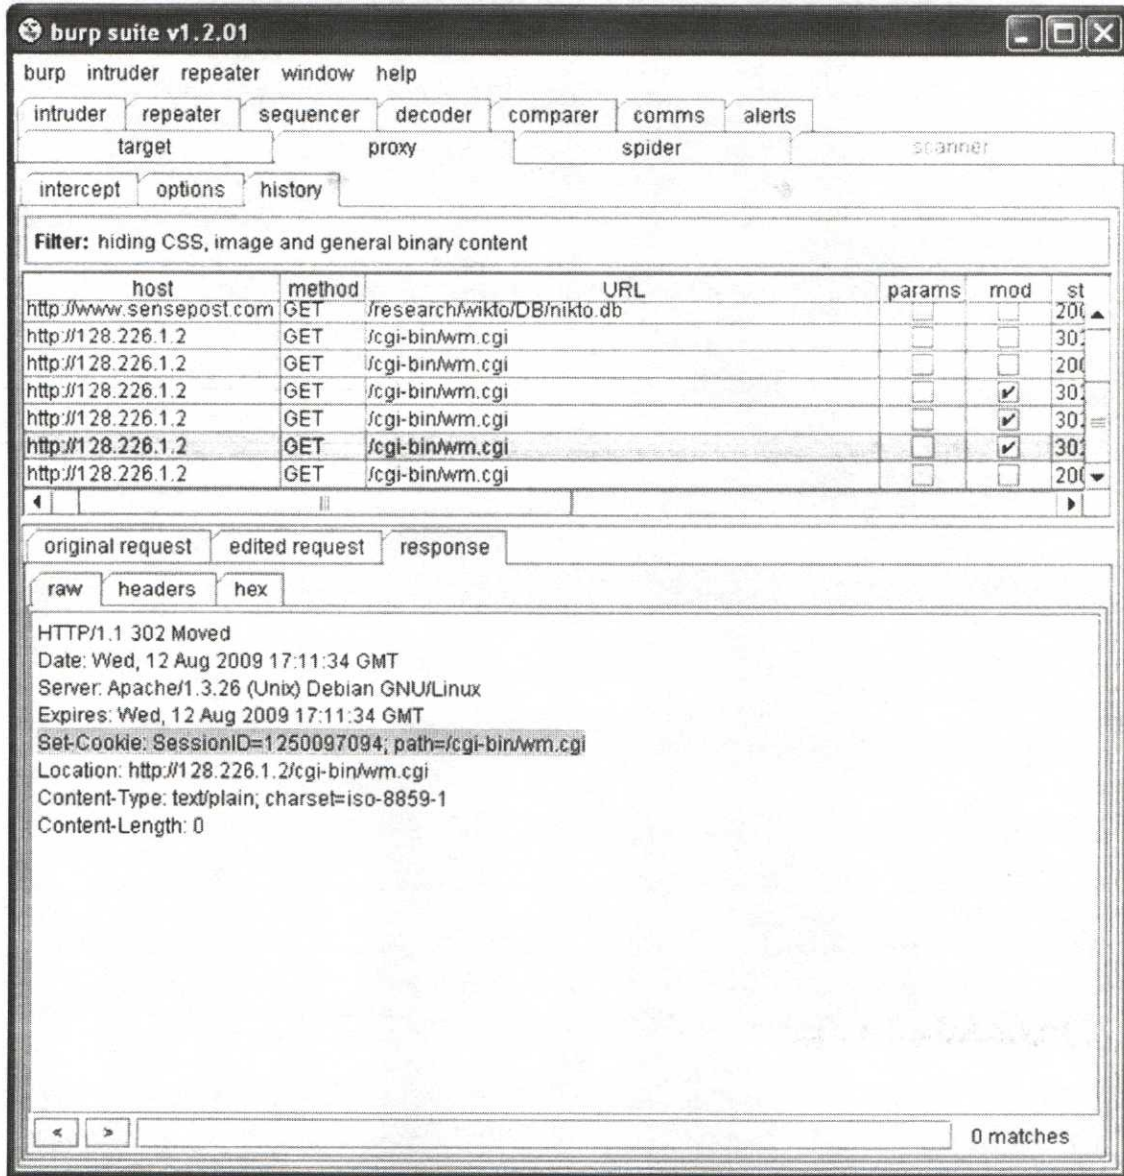
You can see that we have highlighted the request to “/cgi-bin/wm.cgi.” *Please highlight this line in your Burp window too. After selecting this conversation, please click the “Request” tab on the bottom of the window.* If you scroll down, you will also see that the “SessionID” cookie has been sent to the application.

To clear the SessionID cookie most easily, please switch to the “Proxy” tab and turn on the “Intercept” option, and then click the Refresh button in your web browser.



After clicking Refresh, you should see a request like the above appear in the Intercept tab. Please highlight, and then delete the “Cookie: SessionID=...” line from the request. After deleting this line, click the “Forward” button.

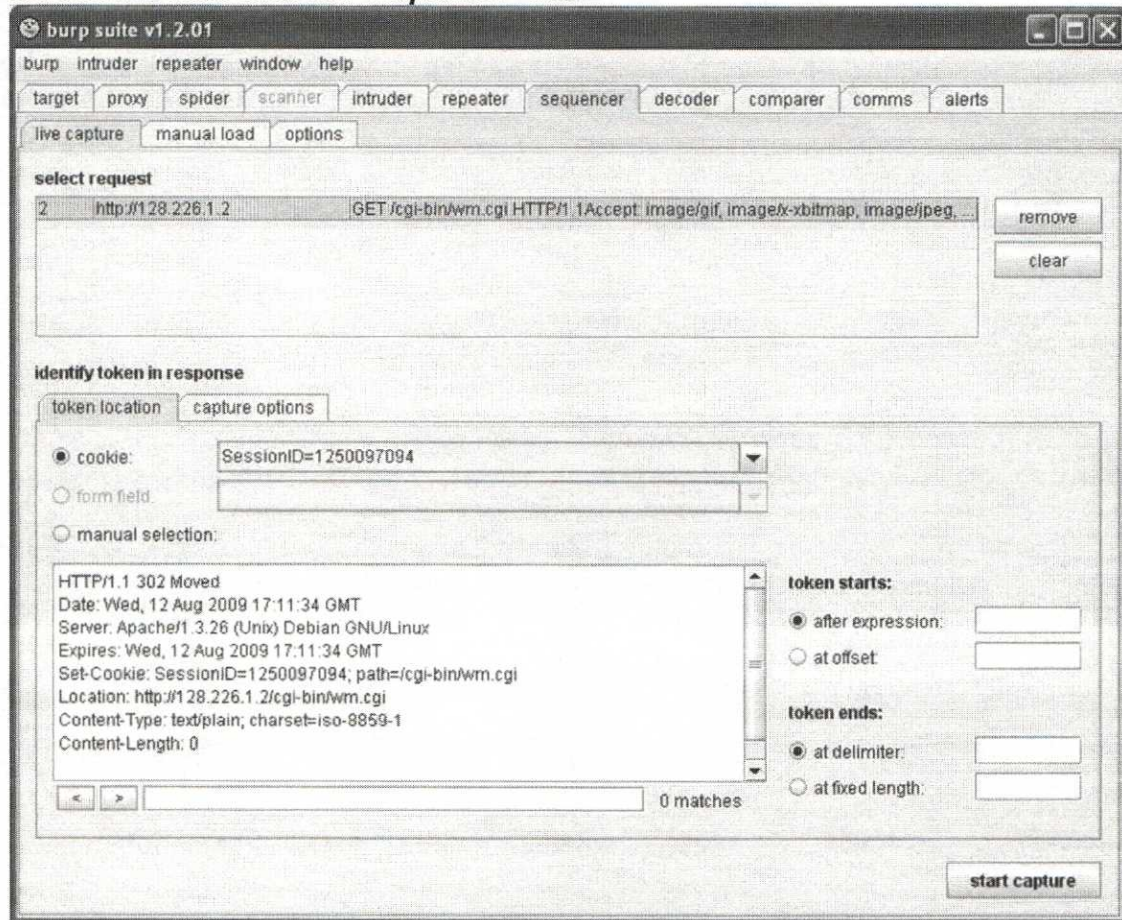
At this point, your web browser might resend the request. Turn off the Intercept option and forward the current request. After this has completed, please click the “History” tab in the “Proxy” section:



You should be able to find that either the last item marked "Modified" (checkmark in the "mod" column) has a "Set-Cookie" header in it. After you find this item, right-click it and choose "Send to Sequencer."

What we have just done is forced the web server to send us a new session ID by removing the existing session ID from the original request. Now that we have a fresh session ID being sent with the "Set-Cookie" header, we can allow the Sequencer to automatically extract it from the request.

The sequencer can actually extract session IDs from anywhere in the request, as you can see below. *Please click the “Sequencer” tab.*



Please click the “start capture” button on the lower right-hand corner.

This step will actually take some time to complete. The Burp Sequencer is now sending request after request in an attempt to harvest large numbers of session IDs from the application. The actual amount of time that this will take completely depends on the amount of memory and the speed of your computer. Working with virtual systems requires patience, of course. Using this tool on real systems takes far less time. Using the purchased version of the tool makes it even faster!

After the sequencer completes, please look through the results in the sequencer window. How do these reports compare to what is done by WebScarab?

When would WebScarab be more useful? When would Burp be more appropriate?

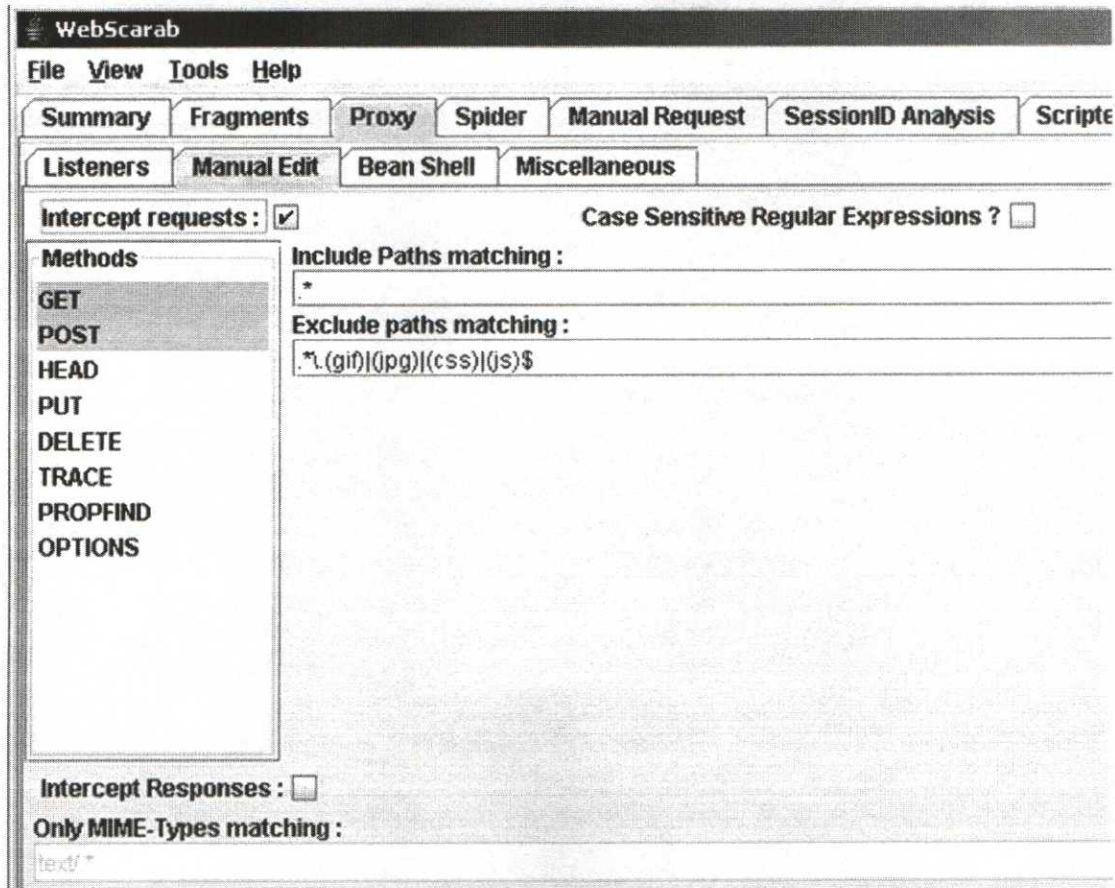
Complete Audit

Purpose: Allow the you the time and test-bed to examine several web applications using the tools and techniques discussed during the day with an organized checklist style approach.

Before the auditing begins, take a look at a few more wrap-up features and questions using WebScarab. Once these have been done, you will be ready to go!

Input Validation and Manipulation

Please select the “Proxy” tab in WebScarab. Notice that a second set of tabs appears beneath the primary tabs. Please select the “Manual Edit” tab. Here is where we can configure WebScarab to allow us to modify requests before they are sent to the web server! Turn on the “Intercept Requests.”



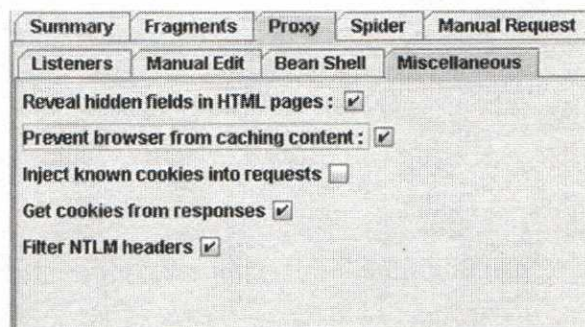
Although it is possible to intercept responses, why is this of less value?

Can you think of something useful that we might want to do that could be accomplished by intercepting responses?

(Hint: Are there things that are sometimes hidden in web pages?)

(Hint: Things that we might want to edit or change before submitting the page?)

Please click the “Miscellaneous” tab. Which options here are valuable and why?



Please set your “Miscellaneous” options as pictured:

Switch to your web browser and click the “Reload” button. You will notice that the web browser seems to get stuck.

What has WebScarab done?

Click the “Accept Edits” button to allow the request to go to the web server. You will see the response come back to WebScarab very quickly. You will need to “Accept Edits” on this one as well.

What happens after you send the first server response?

Hidden Content

Please take note that on the home page of our buggy bank that we have two account numbers and PINs. We’ll note them here now because we’re going to want to use them later:

Account # 1234567890123750 (PIN 1234)

Account # 1234567890123850 (PIN 4321)

Click the “Login” link.

Please look at the conversations related to this request using the “Show Conversation” option.

What information can you find in the Server Data window that reveals the possible existence of a vulnerability on the host?

The vulnerability will essentially be a confidentiality issue. If you could read the file in question from the server, why is this a serious confidentiality issue (Confidentiality as in CIA...)?

Can you actually download this file?

What kind of web server is running on the remote host?

Why is this information valuable to an attacker?

Why would this information be important to a customer?

Cache Prevention

Please log in to one of the accounts and obtain balance information. Do you see any cache prevention efforts in the headers returned from the server?

Why is this a security issue? Who is MOST directly impacted?

Please suggest a single modification to the page that should prevent the page from being cached:

Session Tracking

How is the login session tracked?

Click the "Transfer Funds" link. Which HTTP method is used for handling forms?

What is the advantage to using a form?

Are there any hidden form elements in use?

This site, Buggy Bank, has quite a host of vulnerabilities (intentionally!). At this point, we have given you a fast but somewhat gentle introduction to what to look for in web applications at a high level. Please remember that as auditors, we should be focused on the “What” of the audit, at least at the beginning. We should never allow the “How” to slow us down when first determining what it is we need to examine or where we should find controls. To perform the “How,” we can leverage the experience of the technical support staff or perform the research necessary to examine the technology in question. In this case, we have listed the major flaws that exist in Buggy Bank. Should you want to spend time working on these at home or if you would like to give the technical security folks a good exercise, the Buggy Bank software is available free from the OWASP group. You can find information about the software and its author at www.owasp.org under “Web Maven.”

Please experiment and explore the web applications provided using WebScarab. As time permits, try out the session ID analysis tools and the other cool features that WebScarab includes. You might also want to try out the Paros tool. Paros includes some basic vulnerability scanning similar to what might be found in NStealth, though it is somewhat behind WebScarab in terms of other features and capabilities.

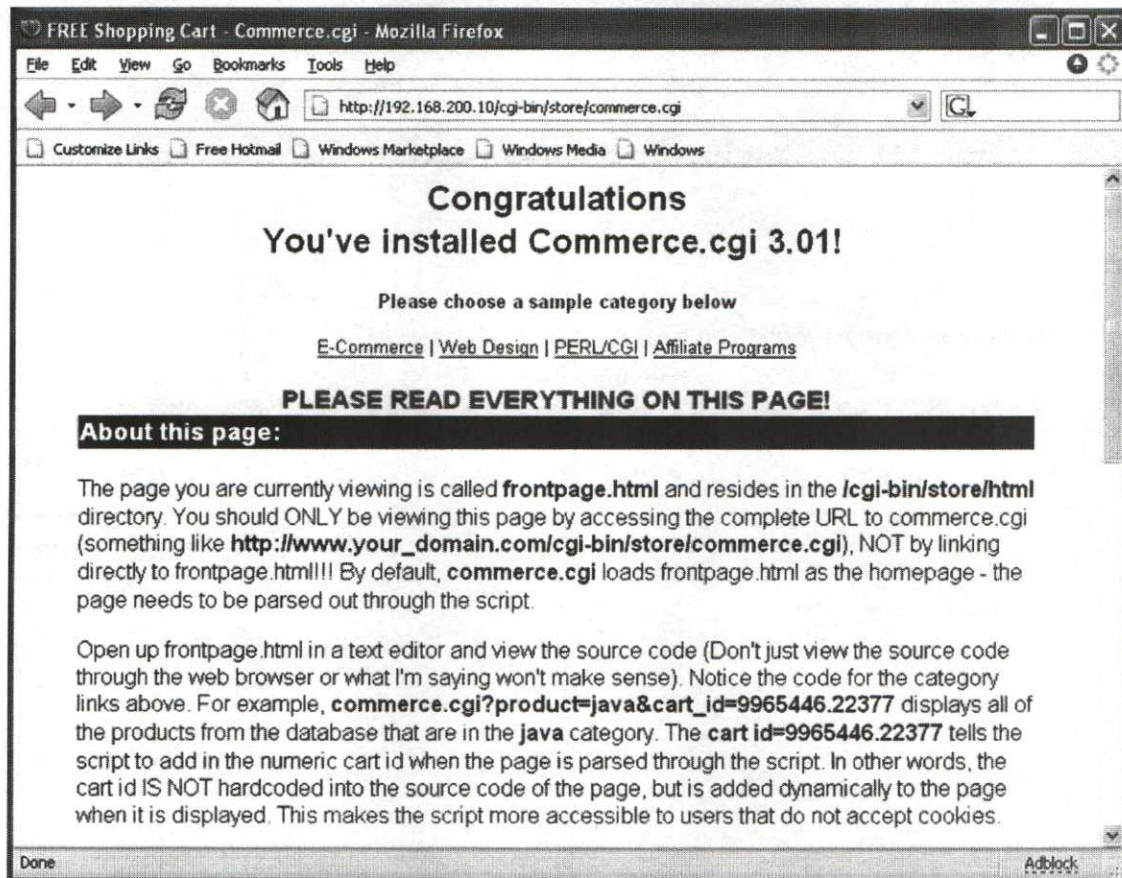
- # - anti-caching techniques are not used on pages displaying sensitive data
- # - predictable session ID in SessionID cookie; session cloning (a.k.a., session hijacking of current sessions)
- # - ODBC/SQL error message when special characters placed in "transaction" parameter
- # - collecting balances for other customers by exploiting a logic flaw
- # - XSS via "transaction" parameter
- # - hidden comment in home page pointing to old CGI source code
- # - user name harvesting
- # - password harvesting via error message for locked accounts
- # - command injection; OS commands embedded after semi-colon (;) in Account cookie (only PING or NETSTAT commands are supported because we don't want anybody hurting themselves :-)
- # - GET method exposes user account number and PIN in URL during login

Web Application Auditing – Typical Homegrown

Purpose: The purpose of this exercise is to give you a measure of “real-world” hands-on experience with a web application that was *not* designed to be insecure. This does not mean that there are no security issues in the application! The application selected is a “real” web application that you can download on the Internet (we included the source code on the course USB as well) and is a good example of what home-grown code looks like when written by someone relatively inexperienced with the vulnerabilities in web applications.

The web application that you will examine is a web-based shopping cart written in Perl. This shopping cart is available “free,” though it is possible to purchase a membership that gets you access to technical support forums, etc.

The web application can be found by starting up the “Unix Web Server” VMware image and connecting to the following URL. Please replace the “xxx.xxx.xxx.xxx” with the address that your Unix Web Server system is running at.
<http://xxx.xxx.xxx.xxx/cgi-bin/store/commerce.cgi>



Why is there a significant level of risk associated with running an “off-the-shelf” web application like this?

Would most organizations claim that they are not using this type of web application?

Even if an organization is using “home-grown” code, how might issues from this example still affect them? (There are several reasons.)

Please use the Web Application Basic Security Checklist that is provided at the back of the workbook to perform a complete audit of this application. At a minimum, please answer these questions. If possible, identify any security issues that can be a problem and rate the severity of these issues.

- How are sessions tracked?
- Are there any anti-caching techniques employed when sensitive information is transmitted or requested?
- Are there any security precautions in place for the transmission of sensitive information?
- Are there any security issues with the authentication scheme in use?
- Is there any attempt to prevent session cloning?

For more thorough and structured testing, please refer to and use the checklist provided in your course book!

This page is provided for your notes while analyzing the web application.

Day 4

Lab Overview

The labs in this manual are intended to reinforce the classroom material and to provide an opportunity for you to gain some hands-on practice with the tools and techniques discussed in the course.

The labs are laid out as follows:

- The labs are designed with Windows in mind because this is the typical Windows operating system that will be found in modern business environments today. 64-bit versions of Windows 7 Professional, Windows 8 Professional, Windows Server 2008, Windows Server 2012, and Windows 10 systems will all work well. If your laptop is loaded with a different version of Windows, the syntax of some commands might vary and not all tools are available.
- When performing labs on your own laptop, keep the following in mind:
 - Under normal circumstances, the labs themselves should not cause harm to your system. However, we cannot predict the interaction of lab tools with every possible laptop hardware and software configuration. **Use these tools on or against your own system at your own risk.**
 - Some labs can actually modify your system's configuration and security settings. If your laptop is loaded with a standard system image, the labs might make undesirable changes to that image. **If you have any concerns about modifying your system configuration, skip the labs that actually require you to make changes or simply do not make the changes when they are called for.**
 - **SANS cannot provide technical support if your laptop does not work or you experience technical problems unrelated to the labs.**
 - **SANS cannot provide you with a copy of Windows if you did not bring a laptop with the correct (laptop requirements) version of Windows installed.**

Each lab contains the following elements:

Objective, describing the goal or purpose of the lab.

Estimated lab time, the estimated amount of time to complete the lab.

Requirements, the tools needed to carry out the lab.

Short description, a high-level overview of the lab that follows.

Detailed description, which provides step-by-step instructions and screen shots.

Optional labs, for advanced students or students who finish early and want to do additional work.

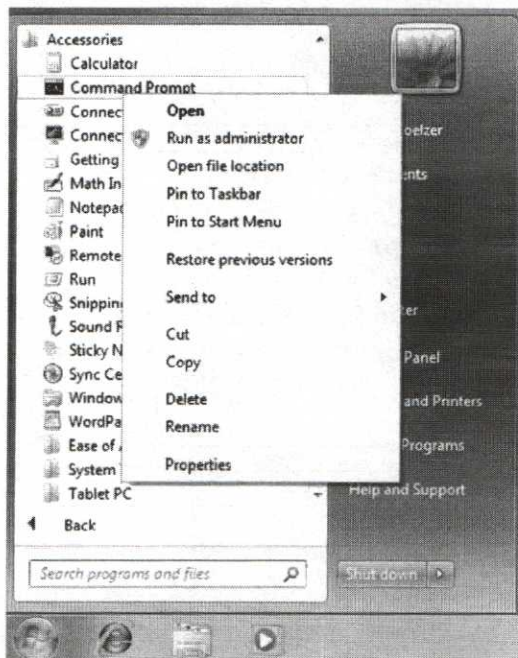
Special Note About Windows Utilities

Windows often includes more than one version of a utility. Generally, this is done to allow access to “older” Windows tools for those users and administrators who might be used to an earlier version of a particular tool.

Two examples are the Windows command prompt and the Windows registry editor.

Command prompt: If you are a long-time Windows user, you remember the Windows command prompt as being launched by the command `command.com`, a 16-bit command line interpreter. Although `command.com` still exists in Windows systems in the NT line (Windows 7, Windows 8, Windows 10, etc.), it has been supplanted by `cmd.exe`, a 32-bit command-line interpreter. Even then, simply attempting to execute “`cmd.exe`” on a 64-bit Windows system, which interpreter is running, the 32-bit or the 64-bit version? It turns out that both are stored on the system. The 64-bit version is in the `c:\windows\system32` directory, whereas the 32-bit version is in the `c:\windows\SysWow64` directory¹.


The command prompt can be accessed by pressing the Windows button at the far left of the Windows task ribbon and either typing “`cmd.exe`” into the search box or locating the “Command Prompt” selection in the Accessories folder.



On modern Windows systems, you will likely want to run the `cmd` prompt as an administrator. Modern versions of Windows have additional security mechanisms. You might find that some tests or commands will not work correctly if you are not running an administrative command prompt even though you might be logged in as an administrator.

To open an administrative command prompt on these systems, begin by clicking the Windows icon that has replaced the “Start” menu. Navigate to “Accessories,” and then right-click the Command Prompt item. In the menu that appears, you will see that the second option is to “Run as administrator” as pictured to the left. Select this option.

¹ This seems very backwards since the 64 bit utilities are in “system32” while 64 bit are in “Wow64.” It’s not very important, but it is good to know.



Connecting to the Cloud

How you connect your system to the lab environment will depend on how you are consuming this class. If you are taking this class via the OnDemand system, for example, you will want to review and follow the directions on the page titled, “OnDemand Cloud Access.” On the other hand, if you are at a live conference, you will need to follow the appropriate directions for the venue that you are attending.

OnDemand Cloud Access

The following instructions will allow you to interface with an example Active Directory server during the Day 4 labs and to the *NetWars: Audit the Flag* experience during Day 6. Before starting, as an OnDemand student, you should have received a digital certificate with your course materials, likely via email. Please locate this certificate now. ***You will be unable to complete these instructions without this certificate.***

We will use an SSL-based VPN for this class. To access it, you must first install OpenVPN. ***In the “Days 1-5” folder on the USB, you will find a folder labeled “VPN.” Open this folder.*** The installer for OpenVPN is in this folder.

Please begin by right-clicking the openVPN-install-2.3.0-1005-i686.exe installer and choosing “Run as Administrator.” This installer requires the installation of a tunneling driver, which will allow the VPN to route traffic between your computer and the VPN network. ***Please accept all defaults.*** When the installation finishes, it will offer to show you the “README” file. You are welcome to peruse the file, but it is perfectly fine to skip this.

The next step is to copy the appropriate configuration files into the location where the VPN service will look for them. To do so, ***please open a Windows file browser and find the OpenVPN installation folder.*** The installation folder should be “C:\Program Files\OpenVPN”.

Within the OpenVPN directory is a directory named “config.” ***Please drag the “VPNConfig” folder from the VPN folder on the USB within the “Days 1-5” folder and drop it into the “config” folder.*** This places the appropriate configuration file into the VPN configuration directory along with the necessary certificates. Finally, ***locate the certificate issued to you by SANS. Copy this certificate into the VPNConfig folder within the “config” folder in the OpenVPN directory.***

The last step required will be to actually run the VPN client. There is no need to connect at the moment, but if you’d like to test it out, here are the next steps: ***Right-click the “OpenVPN GUI” icon that was placed on your desktop during the installation of openVPN. After right-clicking, choose the “Run as Administrator” option.*** It is absolutely critical that you start the VPN client as an administrator. If you fail to do so, you will seem to successfully connect to the VPN, but none of your data will route to the VPN. The reason is that adjusting the route table ***requires*** that you are an administrator.

That’s it! No need to worry about this at the moment. We will put this into use in class as the week progresses.

Live Conference Cloud Access

The following instructions allow you to interface with an example Active Directory server during the Day 4 labs and to the *NetWars: Audit the Flag* experience during Day 6. Before starting, your instructor will provide each you with a copy of an authentication certificate that has been generated for your class. ***You will be unable to complete these instructions without this certificate.***

We will use an SSL-based VPN for this class. To access it, you must first install OpenVPN. ***In the “Days 1-5” folder on the USB, you will find a folder labeled “VPN.” Open this folder.*** The installer for OpenVPN is in this folder.

Please begin by right-clicking the OpenVPN installer and choosing “Run as Administrator.” This installer will require the installation of a tunneling driver, which will allow the VPN to route traffic between your computer and the VPN network. ***Please accept all defaults.*** When the installation finishes, it will offer to show you the “README” file. You are welcome to peruse the file, but it is perfectly fine to skip this.

A pre-configured VPN configuration file has been prepared for your class. You can obtain this file from <http://auditcasts.com/vpn.zip>

After downloading the zip file containing the configuration file, either tell the web browser to open it or double click on the zip file in your downloads directory. Within the archive is one file. Please right-click on this file and select “Copy.”

The next step is to place the configuration file into the location where the VPN service will look for it. To do so, ***please open a Windows file browser and find the OpenVPN installation folder.*** The installation folder should be “C:\Program Files\OpenVPN”.

Within the OpenVPN directory is a directory named “config.” ***Please open the “config” folder and then right-click inside of it and choose “Paste.”*** This places the appropriate configuration file into the VPN configuration directory

The last step required will be to actually run the VPN client. There is no need to connect at the moment, but if you’d like to test it out, here are the next steps: ***Right-click the “OpenVPN GUI” icon that was placed on your desktop during the installation of openVPN. After right-clicking, choose the “Run as Administrator” option.*** It is absolutely critical that you start the VPN client as an administrator. If you fail to do so, you will seem to successfully connect to the VPN, but none of your data will route to the VPN. The reason is that adjusting the route table ***requires*** that you are an administrator.

That’s it! No need to worry about this at the moment. We will use this in class as the week progresses.

WMIC Exercise

This lab is strictly an information-gathering exercise; no changes will be made to the system.

Objective: Experiment with the Windows Management Instrumentation Console tool to understand its basic functioning and discover the types of information that can be retrieved from a system.

Estimated lab time: 15 minutes

Requirements:

- Laptop running Windows XP, Vista, 7, 8, or 10

Lab – Short Description:

A short description of the lab steps is included for those students who are already familiar with WMIC and are capable of working through the lab without detailed instructions. Those students who are less familiar with the subject matter should use the Detailed Description below, which includes step-by-step instructions.

1. Using the Windows command line, ensure that WMIC is properly initialized.
2. Use WMIC to gather the following information about your local host:
 - Items that will run at startup
 - List of all services
 - Current process list
 - Physical drive configuration
 - Logical disk configuration
 - List of installed software

Lab – Long Description:

The long description of the lab will walk you through gathering each of the items required step by step. Along the way, we will include some discussion of why the information is useful to us as auditors and interesting findings that you might see in a live environment.

1. Please open a Windows command prompt as described in the “Special Note About Windows Utilities.”
2. In the command prompt window, please enter “wmic” and press Enter. This will give you an interactive session with the WMIC command-line interpreter (CLI).
3. At the CLI prompt, please type “/?.” The “/?” command will

```
Administrator: Command Prompt - wmic
Microsoft Windows [Version 6.1.7600]
Copyright (c) 2009 Microsoft Corporation. All rights reserved.

C:\Windows\system32>wmic
wmic:root\cli>_
```

display the “Help” page. If you were running this tool non-interactively, you could obtain this information by typing “wmic /?” at the command prompt. This is why the command to obtain help is “/?”.

```

Administrator: Command Prompt - wmic
wmic:root\cli>/?

[global switches] <command>

The following global switches are available:
/namespace Path for the namespace the alias operate against.
/role Path for the role containing the alias definitions.
/node Servers the alias will operate against.
/implevel Client impersonation level.
/authlevel Client authentication level.
/locale Language id the client should use.
/privileges Enable or disable all privileges.
/trace Outputs debugging information to stderr.
/record Logs all input commands and output.
/interactive Sets or resets the interactive mode.
/failfast Sets or resets the FailFast mode.
/user User to be used during the session.
/password Password to be used for session login.
/output Specifies the mode for output redirection.
/append Specifies the mode for output redirection.
/aggregate Sets or resets aggregate mode.
/authority Specifies the <authority type> for the connection.
/?[:<BRIEF!FULL>] Usage information.

For more information on a specific global switch, type: switch-name /?

The following alias/es are available in the current role:
ALIAS - Access to the aliases available on the local system
BASEBOARD - Base board (also known as a motherboard or system board)
BIOS - Basic input/output services (BIOS) management.
BOOTCONFIG - Boot configuration management.
CDROM - CD-ROM management.
COMPUTERSYSTEM - Computer system management.
CPU - CPU management.
CSPRODUCT - Computer system product information from SMBIOS.
DATAFILE - DataFile Management.
DCOMAPP - DCOM Application management.
DESKTOP - User's Desktop management.
DESKTOPMONITOR - Desktop Monitor management.
DEVICEMEMORYADDRESS - Device memory addresses management.
DISKDRIVE - Physical disk drive management.
DISKQUOTA - Disk space usage for NTFS volumes.
Press any key to continue, or press the ESCAPE key to stop
  
```

(For your convenience, we have included a list of the most commonly available WMIC aliases at the back of this workbook starting on page 235.)

- To obtain information about any of the items listed, you simply need to enter the name of the WMI facility that you would like to access followed by the “LIST” command. The first item that we would like to obtain is a list of all of the services that will run at startup. To do this, please type “startup list” at the CLI.

```

wmic:root\cli>startup list
Caption Command Description
Sidebar %ProgramFiles%\Windows Sidebar\Sidebar.exe /autoRun Sidebar
Sidebar %ProgramFiles%\Windows Sidebar\Sidebar.exe /autoRun Sidebar
VMware Tools "C:\Program Files\VMware\VMware Tools\VMwareTray.exe" VMware
VMware User Process "C:\Program Files\VMware\VMware Tools\VMwareUser.exe" VMware

wmic:root\cli>
  
```

(Important! You might find that your attempt to enumerate the startup items fails. Some antivirus products intercept attempts to view the startup items. If you want to correct this, you will need to temporarily disable your antivirus tool.

As an alternative, use the “/?” option to identify some other item to enumerate in your system.)

5. Please use the “/?” command to find the WMI facility names to obtain each of the following items. After you find the WMI name, use “LIST” to retrieve the data:
 - List of services
 - Physical drive configuration
 - Current process list
 - Logical drive configuration
 - List of installed software
6. At the CLI prompt, please type “QUIT.” This will exit the WMIC CLI. At the command prompt, please type the following:
`wmic startup list`
(If your antivirus prevented the use of the startup list option, please use any other option that was successful.)
7. As you can see, this retrieves the same information and requires only a single command to be run at the command prompt. Compare the information shown with the “list” command to what is displayed with “list brief”:
`wmic startup list brief`

```
C:\Windows\system32>wmic startup list brief
Caption      Command                                     User
Sidebar      %ProgramFiles%\Windows Sidebar\Sidebar.exe /autoRun  NT AUTHORITY\LOCAL SERVICE
Sidebar      %ProgramFiles%\Windows Sidebar\Sidebar.exe /autoRun  NT AUTHORITY\NETWORK SERVICE
VMware Tools "C:\Program Files\VMware\VMware Tools\VMwareTray.exe"  Public
VMware User Process "C:\Program Files\VMware\VMware Tools\VMwareUser.exe"  Public

C:\Windows\system32>
```

8. Please try retrieving all of the items from step 5 from the command prompt rather than the CLI.

Alternative Approach – Remote Server

Throughout the day today, remember that there is a Windows domain controller available for you to interact with. To connect to the VPN where it is connected, you must first follow the directions on page 20. After you have successfully connected, you can use WMIC to complete the above exercise against this system.

An example command line to connect to this remote system is:

```
wmic /node:507dc.enclaveforensics.com /user:auditor /password:Password1 service list
```

Basic Scripting Exercise

This lab is strictly an information-gathering exercise; no changes will be made to the system.

Objective: Learn basic scripting techniques. Examine an automated means of finding data in automatically generated text files.

Estimated lab time: 15 minutes

Requirements:

- Laptop running Windows 7, Server 2008, 8, 10, or Server 2012

Lab – Short Description:

A short description of the lab steps is included for those students who are already familiar with basic PowerShell scripting and are capable of working through the lab without detailed instructions. Those students who are less familiar with the subject matter should use the Detailed Description below, which includes step-by-step instructions.

Create a PowerShell script that will compare the list of services output by the WMIC Services list to a previously obtained list of services. If differences are detected, your script should alert the operator.

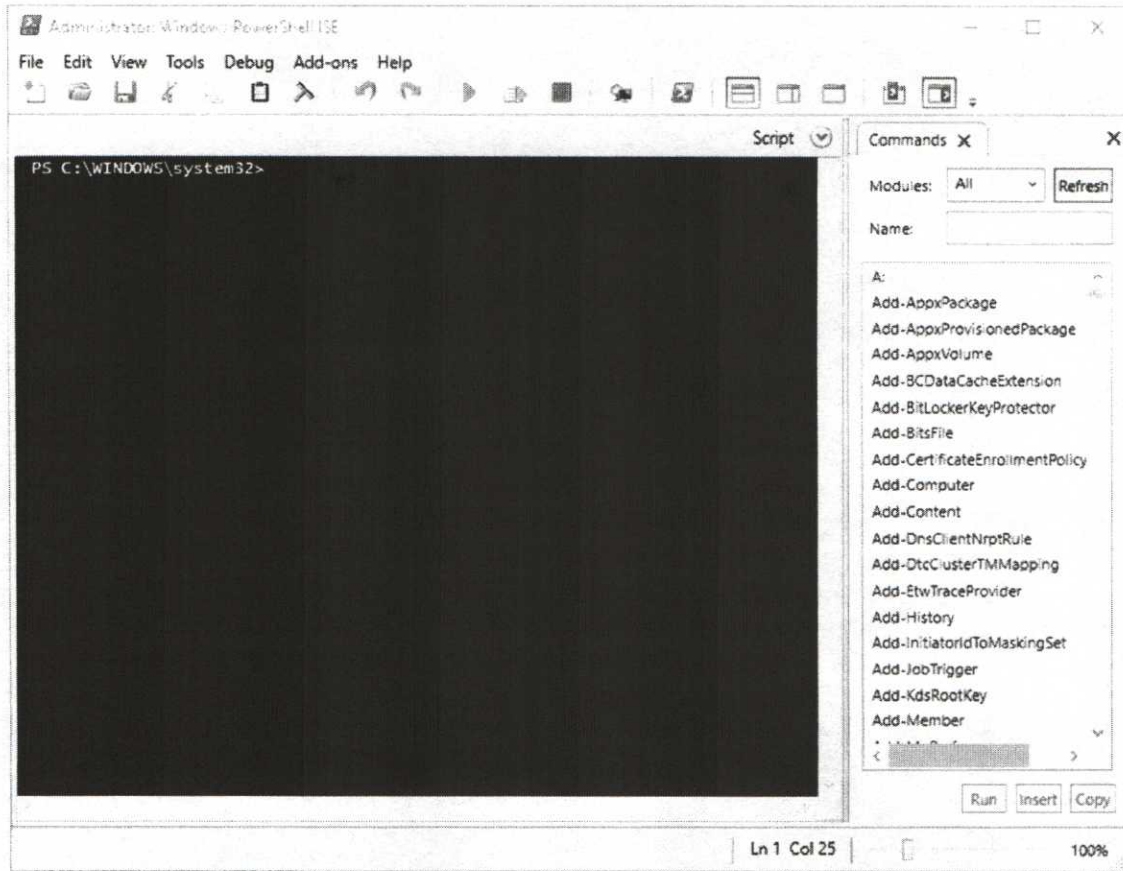
Lab – Long Description:

This lab lays out some basic principles of scripting in a logical order. Rather than focus all of your energy on learning all of the ins and outs of scripting, it is much more important to learn to recognize “recipes” that can be easily modified by you to implement automatic testing and compliance monitoring in your environment.


1. To begin, please use your Start menu or Windows search bar to locate “PowerShell ISE.” To execute it, please right-click it and run it as an administrator.

***Special Note for Windows 8:**

Under Windows 8 you will not be able to locate the ISE directly. To get it started, simply search for “Powershell.” This will allow you to execute a Powershell prompt. ***Right-click on the Powershell command and run it as an Administrator.*** Now that you have Powershell running, you can get the ISE to start up. ***Please type ‘ise’ at the Powershell command prompt and the ISE will open.***



Notice that in the main window (to the left), you now have what appears to be a command prompt¹, and on the right-hand side you have a quick reference help window for PowerShell scripting commandlets. There is also a button labeled “Script” on the top right edge of the command prompt. Click this button to reveal a script-editing window.



2. Create a folder on your hard drive named “Scripts.” We recommend that you create this folder in the root of the C drive so that it will be very easy for you to find when you test your script.
3. Use the “CD” command to change to the “Scripts” directory that you have created (CD\Scripts).

¹ In fact, it is a command prompt, but it’s not the command prompt that you might be accustomed to. All of the same commands will work as expected, but this is actually a PowerShell command prompt.

```

PS C:\WINDOWS\system32> cd \
PS C:\> mkdir Scripts

Directory: C:\

Mode                LastWriteTime         Length Name
----                -
d-----            1/19/2016   5:45 AM             Scripts

PS C:\> cd Scripts
PS C:\Scripts> |

```

4. *At the Powershell command prompt, please enter the following:*

```

PS C:\Users\dhoel> Set-ExecutionPolicy RemoteSigned

Execution Policy Change
The execution policy helps protect you from scripts that you do not trust. Changing the execution policy might
expose you to the security risks described in the about_Execution_Policies help topic at
http://go.microsoft.com/fwlink/?LinkID=135170. Do you want to change the execution policy?
[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help (default is "N"):

```

5. **Select the 'A' option.** The ExecutionPolicy determines which scripts it is permissible for your Powershell installation to run. The execution policy will be configured differently depending on the version of Windows that you have and what settings your administrators have set as defaults. Selecting 'RemoteSigned' will allow us to run the script that we write locally but would require remote scripts to be digitally signed. This is a good balance that allows us to script easily will still enforcing some controls on untrusted scripts.
6. Within the open script editing window, please enter the following line:

```
# A First Powershell Script
```

The leading hash mark indicates that this is a comment. In fact, you can use the double slashes anywhere in your script; everything on the line following them will be treated as a comment and will not be executed.

7. Our script must obtain the filename to compare the list of services to from the command line. To do this, we will take advantage of the built-in access to parameters that are passed using a dash. (For example, "script -baseline base.txt" calls a script named "script" with a parameter named "baseline" that has a value of *base.txt*.) Let's document this fact and add the code that lets us access the parameter:

```
# User must provide a "-base" parameter that specifies
a baseline file.
param($base="baseline.txt")
```

8. Our script must also generate a list of services that will be used in the comparison. To generate this, we can simply use the WMIC tool¹ to generate the list and store it into a temporary file. Please add the following line to your Notepad window:

```
wmic service list > TempFile
```

9. Now that the current services list is in the temporary file, we need to compare the two files. PowerShell has a very powerful comparison commandlet built into it named "Compare-Object." You can also access this same commandlet using an alias, "diff." Because PowerShell uses objects for everything that it does, we need to turn the baseline and observed files into objects that we can compare. To do this, we will load them into objects "\$a" and "\$b" using the "Get-Content" commandlet. Please add the following to your script:

```
$a = Get-Content $base
$b = Get-Content .\TempFile
```

We're now ready to actually compare the files and act if there is a difference.

10. When we use the "Compare-Object" or "diff" commandlet, we can interrogate it. For example, we can ask² it how many differences there were.

```
if ((diff $a $b).count -gt 0)
{
    echo Changes found!
} else
{
    echo No differences found
}
```

11. Pull down the File menu and select the "Save" option. In the window that appears, browse to your "Scripts" directory. Change the filename to exactly what appears on the next line, including the quotation marks!

```
"script.ps1"
```

12. Before we proceed, let's ensure that there will be a potential finding. Let's start by making sure that the Universal Plug-And-Play service³ is enabled. To do so,

¹ There is a much, much better way to do this using a PowerShell commandlet called "Get-WmiObject." For now, we take a very linear approach and try to limit the amount of PowerShell that we need to learn.

² Again, there are far more things that we can do here, but we're keeping things simple for now. We encourage you to experiment and explore!

³ The service itself is not terribly important. If, for some reason, this service does not exist, examine your list of services and identify something that seems to be innocuous, like the Xbox Service, to use in its place.

please type within the PowerShell command window:

```
net start upnphost
```

13. Next, to test the script, we must create a baseline file that contains a list of services. To do this, please enter the following command in the PowerShell command window:

```
wmic service list > baseline.txt
```

14. The name of the file is completely arbitrary. We have selected the above filename just so that it is obvious what we are doing. Next, we can run the script with “-base baseline_services” as an argument to see whether everything is working properly. Please run the following command in the command prompt window:

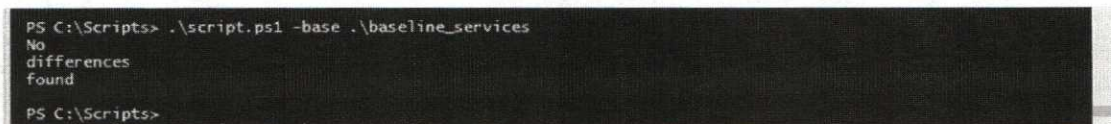
```
.\script.ps1 -base baseline.txt
```

15. The output of the command in the previous step should be, “No differences found.” Next, we want to make sure that the script will detect an actual change. To do this, let’s turn off a service that is typically running. We’ll turn off the Universal Plug-And-Play Host. To do this, please type the following:

```
net stop upnphost
```

16. After a few seconds, the command will complete. Once this is done, please rerun the command from number 12 above. You should now see that differences have been found. Our script is a success!

At this point, we have a very basic “recipe” that can be applied to any task that can be run at the command line and that will produce text output. As we work throughout the day, consider taking the commands demonstrated and trying to re-implement this script with those commands most useful to you.



```
PS C:\Scripts> .\script.ps1 -base .\baseline_services
No
differences
found
PS C:\Scripts>
```

```
1 # A First Powershell Script
2 # User must provide a "-base" parameter that specifies a baseline file.
3 wmic service list > TempFile
4 $a = Get-Content .\baseline_services
5 $b = Get-Content .\TempFile
6 if((Compare-Object $a $b).count -gt 0)
7 {
8     echo Changes found!
9 } else {
10     echo No differences found
11 }
```

Basic System Information/Open Ports and Running Services

This lab is strictly an information-gathering exercise; no changes will be made to the system.

Objective: Use various command-line tools to obtain Service Pack/patch level information about a local or remote system.

Estimated lab time: 15 minutes

Requirements:

- Laptop running Windows XP, Vista, 7, 8, or 10

Lab – Short Description:

A short description of the lab steps is included for those students who are already familiar with Windows XP and related tools and are capable of working through the lab without detailed instructions. Those students who are less familiar with the subject matter should use the Detailed Description below, which includes step-by-step instructions.

1. Install System Internals' PSTools.
2. Use PSTools to obtain basic information about the target system.

Lab – Long Description:

1. Create a new directory called "Audit" to store the output of your audit tools:
mkdir audit
2. Enter the directory listing (dir) command to confirm that you created the directory

Part 1 – Install System Internals' PSTool

1. Locate the folder containing System Internals' PSTools. (If you followed the pre-class setup instructions, it should be located in C:\Tools\Windows Tools\PSTools.)
2. Select all of the files in the PSTools folder and copy them into the C:\Tools directory.

Part 2 – Use PSTools to Obtain Basic Information about the System

Information such as the operating system and version, hostname, system uptime, and hardware configuration will help you identify standard characteristics about the target host. Any changes to this information over time could indicate that the host has been updated or reinstalled, or that a different host is now using the IP address.

1. Open a command prompt window (Start → Run → cmd.exe).
2. Change to your tools directory (CD \tools).

3. Run the following command to view a list of options for the `psinfo.exe` tool:

```
psinfo /?
```

4. Note that `psinfo` allows you to specify a remote hostname or IP address, as well as an account to use (username and password) to connect to the remote machine. For purposes of this lab, you can run `psinfo` against your own system or the instructor's target system. Run `psinfo` with the `-s` switch to also show installed software.

To run the tool against your local system:

```
psinfo -s > c:\audit\psinfo.txt
```

This tool can also be run against a remote system if you have administrative credentials for the remote system:

```
psinfo -s \\remote_system > c:\audit\psinfo_remote.txt
```

To use `psinfo` against the 507dc system, you must first authenticate to the system as follows:

```
net use \\507dc.enclaveforensics.com /user:auditor
```

When you press Enter, you should be prompted for the password. The password is "Password1." At this point, you should be able to execute `psinfo` against this remote system.

5. View the file `psinfo.txt` that you just created. You should see something similar to the following.

```

PsInfo v1.31 - local and remote system information viewer
Copyright (C) 2001-2002 Mark Russinovich
Sysinternals - www.sysinternals.com

Querying information for ...
Uptime:                0 days, 0 hours, 3 minutes, 20 seconds
Kernel version:        Microsoft windows 2000, Uniprocessor Free
Product type:          Advanced Server
Product version:       5.0
Service pack:          3
Kernel build number:   2195
Registered organization:
Registered owner:      Jennifer Kolde
Install date:          10/6/2002, 12:36:02 PM
Expiration date:       2/3/2003, 12:36:02 PM
IE version:             6.0000
System root:           C:\WINNT
Processors:            1
Processor speed:       900 MHZ
Processor type:        Intel Pentium III
Physical memory:       254 MB

Volume Type            Format            Label            Size            Free            Free
-----
A: Removable
C: Fixed               NTFS             win2K_srv        4.0 GB          2.3 GB          57%
D: Fixed               NTFS             winXP_Pro        5.0 GB          2.2 GB          44%
E: Fixed
F: Fixed               NTFS             Data              13.9 GB         10.9 GB         78%
G: CD-ROM              CDFS             DOM34ENUC2       569.5 MB        0%

Applications:
Adaptec Easy CD Creator 4
Adobe Acrobat 5.0 5.0
CIS windows NT/2000 Security Scoring Tool 2.1.3
Cybersafe Log Analyst

```

Part 3 – Identify Open Ports

When analyzing the security of a system, it is important to identify all of the ports that are listening on the network. In particular, we are interested in obtaining a list of open network ports, and then comparing that to the ports that are visible from the local system.

1. Please start up the External virtual machine that was used on Day 2.
The Unix Web Server VM system has Nmap installed. This will give us the opportunity to scan the system remotely without having to use a separate laptop. In a real environment, we would, of course, use an actual external machine to perform the scan. Be aware that when scanning Windows systems, you will either have to scan from the local subnet (which allows Nmap to perform an ARP ping to discover the host) or create an exception in the Windows firewall to permit the scan to occur. In this case, no exception is needed because the virtual machine is local to the Windows host.)
2. After the virtual machine starts, please log in using the credentials *auditor* with a password of *Password1*.

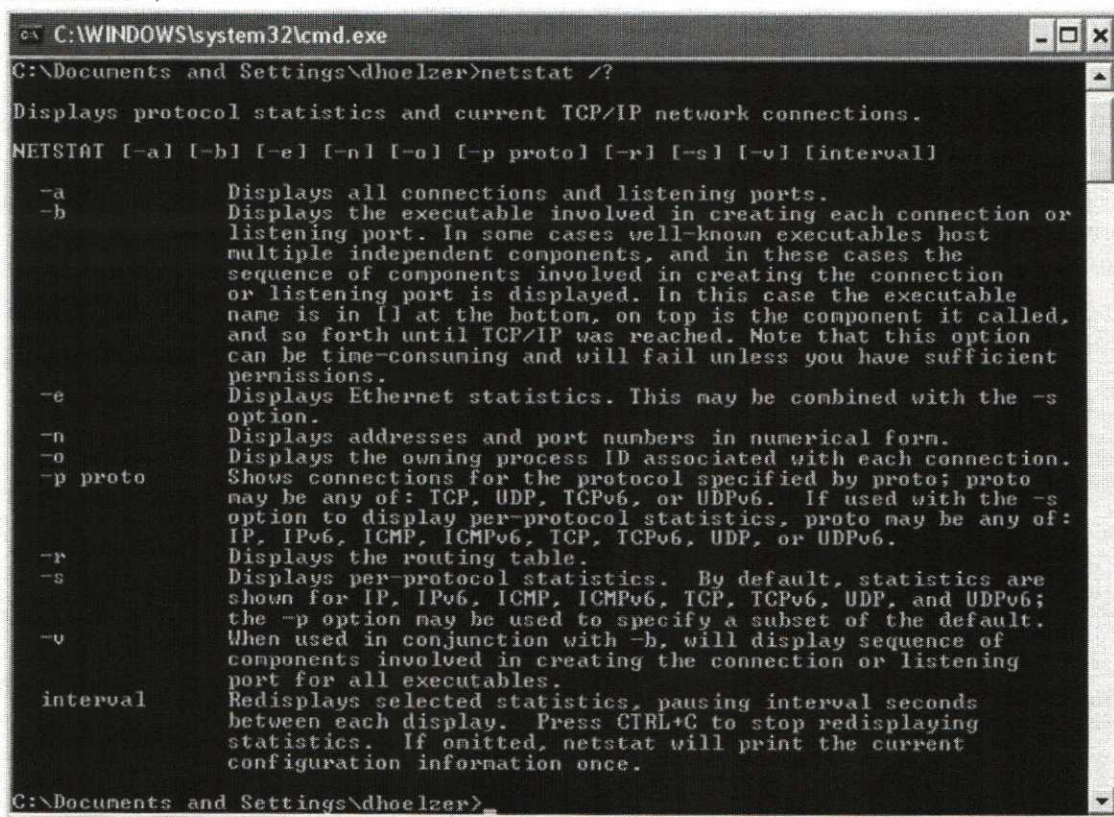
- Please run an Nmap scan of your Windows host. When running the scan, scan all TCP ports with Version and OS fingerprinting enabled.
Immediately above the login banner on the virtual host you will see the IP address that has been assigned to your Unix Web Server system. Your Windows host is on the same subnet; it will be assigned the "1" host address. For example, if your Unix Web Server system reports an address of 192.168.175.129, then your Windows computer will be at the 192.168.175.1 address. The Nmap options that are required for this task are "-p 1-65535 -O."

Part 4 – Use netstat to Identify Services Listening on Each Port

A list of port numbers gives you some idea of the services running on the host. Port 139 is commonly the Microsoft NetBIOS session service, port 389 is commonly LDAP, and so on. But the fact that a particular port **number** is listed is not a guarantee that the service commonly associated with that port is **really** what is listening on that port. Any service can be run on any port; so port 389 could be LDAP, or could be a backdoor Trojan trying to **look** like LDAP. To really determine what process is listening on a given port, you need some method to associate the port with a process.

- Open a command prompt window (Start → Run → cmd.exe).
- Run the following command to view a list of options for the netstat tool:

```
netstat /?
```



```
C:\WINDOWS\system32\cmd.exe
C:\Documents and Settings\dhoelzer>netstat /?
Displays protocol statistics and current TCP/IP network connections.
NETSTAT [-a] [-b] [-e] [-n] [-o] [-p proto] [-r] [-s] [-v] [interval]

-a           Displays all connections and listening ports.
-b           Displays the executable involved in creating each connection or
           listening port. In some cases well-known executables host
           multiple independent components, and in these cases the
           sequence of components involved in creating the connection
           or listening port is displayed. In this case the executable
           name is in [] at the bottom, on top is the component it called,
           and so forth until TCP/IP was reached. Note that this option
           can be time-consuming and will fail unless you have sufficient
           permissions.
-e           Displays Ethernet statistics. This may be combined with the -s
           option.
-n           Displays addresses and port numbers in numerical form.
-o           Displays the owning process ID associated with each connection.
-p proto     Shows connections for the protocol specified by proto; proto
           may be any of: TCP, UDP, TCPv6, or UDPv6. If used with the -s
           option to display per-protocol statistics, proto may be any of:
           IP, IPv6, ICMP, ICMPv6, TCP, TCPv6, UDP, or UDPv6.
-r           Displays the routing table.
-s           Displays per-protocol statistics. By default, statistics are
           shown for IP, IPv6, ICMP, ICMPv6, TCP, TCPv6, UDP, and UDPv6;
           the -p option may be used to specify a subset of the default.
-v           When used in conjunction with -b, will display sequence of
           components involved in creating the connection or listening
           port for all executables.
interval    Redisplays selected statistics, pausing interval seconds
           between each display. Press CTRL+C to stop redisplaying
           statistics. If omitted, netstat will print the current
           configuration information once.

C:\Documents and Settings\dhoelzer>
```

3. Run the `netstat` command to identify processes listening on specific ports on **your local system** (`netstat` cannot be run against a remote machine).

```
netstat -ano > c:\audit\netstat.txt
```

4. View the file `netstat.txt` that you just created.

Question: Can you identify all running executables? How can you use the Task list to assist in this? Are they legitimate? Are there services that should be turned off? Are there services that you cannot identify? When compared to the list of ports found in step 3, are there any services available externally that cannot be accounted for locally with Netstat?

(Answers will vary depending on system configuration. You should be able to identify and confirm the validity of any services running on your host. If there are things you can't identify, you should research them and disable them if they are not necessary.)

Note: *If no ports and processes appear, it could be due to one of two issues: Either your network interface is disabled, or the only processes running are built in to the operating system (i.e., are not part of an independent process).*

5. In Windows 2000, Windows XP, 2003, Windows Vista, Windows 7, 2008, and higher, several services may be started under the generic process "`svchost.exe`". To identify the specific service(s) associated with the `svchost` process, you will need to do some further investigation.

```
tasklist /svc > c:\audit\tasklist.txt
```

Your output should look similar to the following. Note multiple instances of `svchost.exe` running various Windows services.

```

tlist.txt - Notepad
File Edit Format Help
 0 System Process
 8 System
188 SMSS.EXE
212 CSRSS.EXE      Title:
236 WINLOGON.EXE  Title: NetDDE Agent
264 SERVICES.EXE Svcs: Alert, Browser, Dhcp, dmserver, DnsCache, Ever
276 LSASS.EXE      Svcs: PolicyAgent, SamSs
392 termisrv.exe  Svcs: TermService
480 svchost.exe    Svcs: RpcSs
540 svchost.exe    Svcs: EventSystem, Irmon, Netman, NtmsSvc, SENS
588 spoolsv.exe   Svcs: Spooler
616 msdtc.exe     Svcs: MSDTC
740 LLSSRV.EXE   Svcs: LicenseService
776 PPPoEService.ex Svcs: PPPoEService
792 regsvc.exe    Svcs: RemoteRegistry
820 mstask.exe    Svcs: Schedule
900 THotkey.exe   Svcs: THOTKEY
932 vsmon.exe     Svcs: vsmon
972 winMgmt.exe   Svcs: winMgmt
988 svchost.exe   Svcs: wuauServ
1180 dfssvc.exe   Svcs: Dfs
1252 svchost.exe Svcs: Tapisrv
1244 explorer.exe Title: Program Manager

```

Additional information about a particular service can be found by obtaining the process ID (PID) of the process you are investigating, and entering the following (including the quotes), substituting the actual PID number for <processID>:

```
tasklist /FI "PID eq <processID>"
```

Additional information on `svchost.exe` can be found in the Microsoft Knowledge Base.

An alternative set of tools for examining running processes and listening ports are available from Sysinternals. These tools, process explorer, TCPView, and TCPVCon are included in the "PSTools" folder within the "Tools" folder. We encourage you to have a look at these tools. They do not fit directly into the exercises for today because we are working to build a set of command-line tools that can be automated for large-scale audits, but they are invaluable for incident response and troubleshooting on Windows hosts.

LDAP/DSQuery

Objective: Provide the learner with hands-on experience extracting data out of an Active Directory through the use of LDAP queries and the DSQuery tool.

Estimated lab time: 20 minutes

Requirements:

- Laptop running Windows XP, Vista, 7, 8, or 10
- Microsoft RSAT or Server Admin Tools pack

Lab – Short Description:

A short description of the lab steps is included for those students who are already familiar with Windows and related tools and are capable of working through the lab without detailed instructions. Those students who are less familiar with the subject matter should use the Detailed Description below, which includes step-by-step instructions.

1. Install the Microsoft AdminPak tools.
2. Use DSQuery to extract the following:
 - a. A list of all fully distinguished usernames from the Active Directory
 - b. The SamAccountName for all users in the Active Directory
 - c. A list of all computers within the Active Directory
 - d. Full details of the first user in the Active Directory
 - e. Users who are not required to have a password within the Active Directory

Lab – Long Description:

Install DSQuery

DSQuery is a free tool that is included by default on all installs of the Microsoft Server operating systems. It is also provided as a free download within the Remote Server Administration Tools (RSAT) for Windows 8 or the Windows Adminpak (intended for Server 2003).

Although we prefer to use the most current version of the toolkit, there have been issues with the Windows 8 RSAT installer. Most notably, it will completely refuse to install on many Windows 8 systems. ☺ In addition, it cannot be used on any version of Windows older than Windows 8.

For this reason, in the lab, we will use the older “AdminPak.” This installer *will* warn you that some components will not work properly. This is *completely expected*. It also has no impact on the lab because the incompatible tools are not tools that we will use.

Please locate the “WindowsServer2003-KB304718-AdministrationToolsPack” installer in the “Windows Tools” folder from the “Tools” folder located in the “Days 1-5” folder on the UBS stick. Remember that you might have already copied this to your computer on Day 1! If you did, you’ll find the “Windows Tools” folder on your C drive in the

“C:\Tools” directory. *Please right-click the installer and run it as an administrator.* This will get the installer started. *You will be warned about potential incompatibilities in some newer versions of Windows. Please ignore this warning and accept all defaults!*

Now that the AdminPak has been installed, DSQuery is available for use. *Please open a Windows command prompt. Verify that DSQuery installed properly by simply executing the “dsquery” command with no options.* You should see a help screen as a result.

Running Queries

Please verify that you are currently connected to the lab VPN before attempting the next steps.

Determine whether you can identify the command-line options that will allow you to extract user data using dsquery.

(Hint: Try running DSQuery with no options. You should see the information that you are looking for near the bottom.)

Now that we know how to get to user data, let’s see how we can specify servers and credentials. *Please execute the following command:*

```
dsquery user /?
```

Using the output from the above command, please identify the command-line options for the following:

1. Which option allows you to specify the target server to query?
2. Which option allows you to specify the username to log in as?
3. Which option allows you to specify the password to use to authenticate?

Remember that in the lab environment, we have a domain, but your system is not a member of that domain. If you were sitting at your desk in your environment, you would be able to run all of these DSQuery commands without ever specifying the server, username, or password. Also, you will not need administrative rights! As stated in class, unless administrators have gone out of their way to secure it, every user in the Active Directory has the ability to run nearly any query that you can imagine.

Let’s start with an easy query. *Please execute a query that will list all of the fully qualified or distinguished names for the users in Active Directory.*

We already know that “dsquery user” will return the fully distinguished names of users. In the lab, we must retarget the DSQuery tool so that it queries the 507DC server. We can do so like so:

```
dsquery user -s 507dc.enclaveforensics.com -u auditor -p  
Password1
```

Does this command actually return all of the users? Why do you say so?

You probably noticed that it says that the tool reached a “Default Limit.” ***Bring up the help options for DSQuery again and see whether you can find an option that will allow you to change the limit.***

The command that we’re looking for is “-limit”. The next question is what number to set the limit to. The easiest solution is to tell the system that you want “No limit.” To specify this, ***please try the command again with this added option: “-limit 0”***

Using “-filter”

Let’s turn up the difficulty a couple of notches. Although the DSQuery “User,” “Group,” and “Computer” options are nice, there are many, many things that we just can’t do easily using these options. DSQuery does, however, provide more of a “raw” interface to the LDAP data. To access the directory this way, we specify a type of “*.”

Please execute the following command:

```
dsquery * /?
```

How can we specify a filter or query string to find data using the “*” ?

If you look through the help that’s provided, you’ll find that there is a “-filter” option. This is what will allow you to specify arbitrary LDAP queries of your own!

Using your course book as a reference, identify the objectClass and objectCategory values necessary to find User objects within the Active Directory.

Your course book has a page titled, “AD/LDAP Object Cheatsheet,” around page 87 or so. If you refer to this page, you can find the category and class that allows you to find a user object.

Using the information that you found in the last step, create an LDAP filter that will allow you to find Users in a DSQuery command.

Remember that we discussed the formatting of LDAP queries in the lecture. This included the “Prefix” notation approach that is used to write the queries. We also saw that there are operators like the &, | and ! that allow us to create conditions. In this case, we want to find objects where the objectClass is User AND the objectCategory is Person. What would this look like?

We might be tempted to say:

```
((objectClass=User) & (objectCategory=Person))
```

Of course, although this is easy for *us* to understand, the LDAP system requires that we write this differently. In prefix notation, we would write it as follows:

```
(&(objectClass=User)(objectCategory=Person))
```

Now that we know what the filter or query string is, please use this to extract users.

```
dsquery * -filter  
"(&(objectClass=User)(objectCategory=Person))" -limit 0 -s  
507dc.enclaveforensics.com -u auditor -p Password1
```

Whew!! Still, we're not quite there yet. *Please use the help for the "*" option to figure out how to specify attributes to display. Once you have found that, run your query again but display only the SAMAccountName attribute.*

```
dsquery * -filter  
"(&(objectClass=User)(objectCategory=Person))" -limit 0 -s  
507dc.enclaveforensics.com -u auditor -p Password1 -attr  
SamAccountName
```

Wow! That was a lot of work! However, we can now leverage what we did so far to make things easier for us. In fact, *see whether you can modify the last query to extract all of the computer names from the Active Directory.*

Really, this requires that you modify only the objectClass and objectCategory values. Can you figure out which values allow you to pull computer information?

Let's switch back to Users for a moment. *Please run a query that will extract the first users from the Active Directory and display all of the available attributes for that object.*

This might seem hard at first, but in fact you already have all of the pieces that you need. Really, only two pieces of the previous User query need be modified!

```
dsquery * -filter  
"(&(objectClass=User)(objectCategory=Person))" -limit 1 -s  
507dc.enclaveforensics.com -u auditor -p Password1 -attr *
```

Notice that we changed the limit to "1" and that the attribute is now listed as "*." This extracts only one record (the first) and displays every populated attribute from that object!

The last task is actually more complicated. What we'd like to do is identify users who are not required to have a password. This requires that we use those really nasty looking bitwise filters and the UserAccountControl attribute within the Active Directory.

Using the information in the appendix on page 240, see whether you can figure out the numeric value necessary to find users with no passwords.

Also on that same page, identify the LDAP Rule Identifier necessary to perform a bitwise AND against a field.

Put the information gathered together to create a DSQuery command that will extract all users who are not required to have passwords and will print the SamAccountName for those users.

Wow, that's a lot! Let's work through this step by step. First, in the appendix, we find that the Password Required field has a value of 32 (0x20 in hexadecimal). Next, we find that the bottom of that page tells us that the bitwise operator for an AND in LDAP is represented by 1.2.840.113556.1.4.803.

Let's put this together:

```
dsquery * -filter  
"(&(objectClass=User) (objectCategory=Person) (UserAccountControl:1.2.840.113556.1.4.803:=32))" -limit 0 -s  
507dc.enclaveforensics.com -u auditor -p Password1 -attr  
samAccountName
```

If we now want to add more criteria to this (for example, making sure that the account is not disabled), we could add additional bitwise filters!

Users, Groups, and Passwords

Objective: Use password cracking/password assessment tools to audit the strength of Windows passwords.

Estimated lab time: 15 minutes

Requirements:

- Laptop running Windows XP, Vista, 7, 8, or 10
- Cain and Abel (<http://www.oxid.it>)

Lab – Short Description:

A short description of the lab steps is included for those students who are already familiar with Windows and related tools, and are capable of working through the lab without detailed instructions. Those students who are less familiar with the subject matter should use the Detailed Description below, which includes step-by-step instructions.

1. Install Cain and Abel password cracking tool.
2. Create a number of sample accounts on Windows with passwords of varying strengths.
3. Locate the hashes from the in-class server.
4. Configure Cain and Abel to extract a copy of the password hashes and attempt to crack the passwords.

Lab – Long Description:

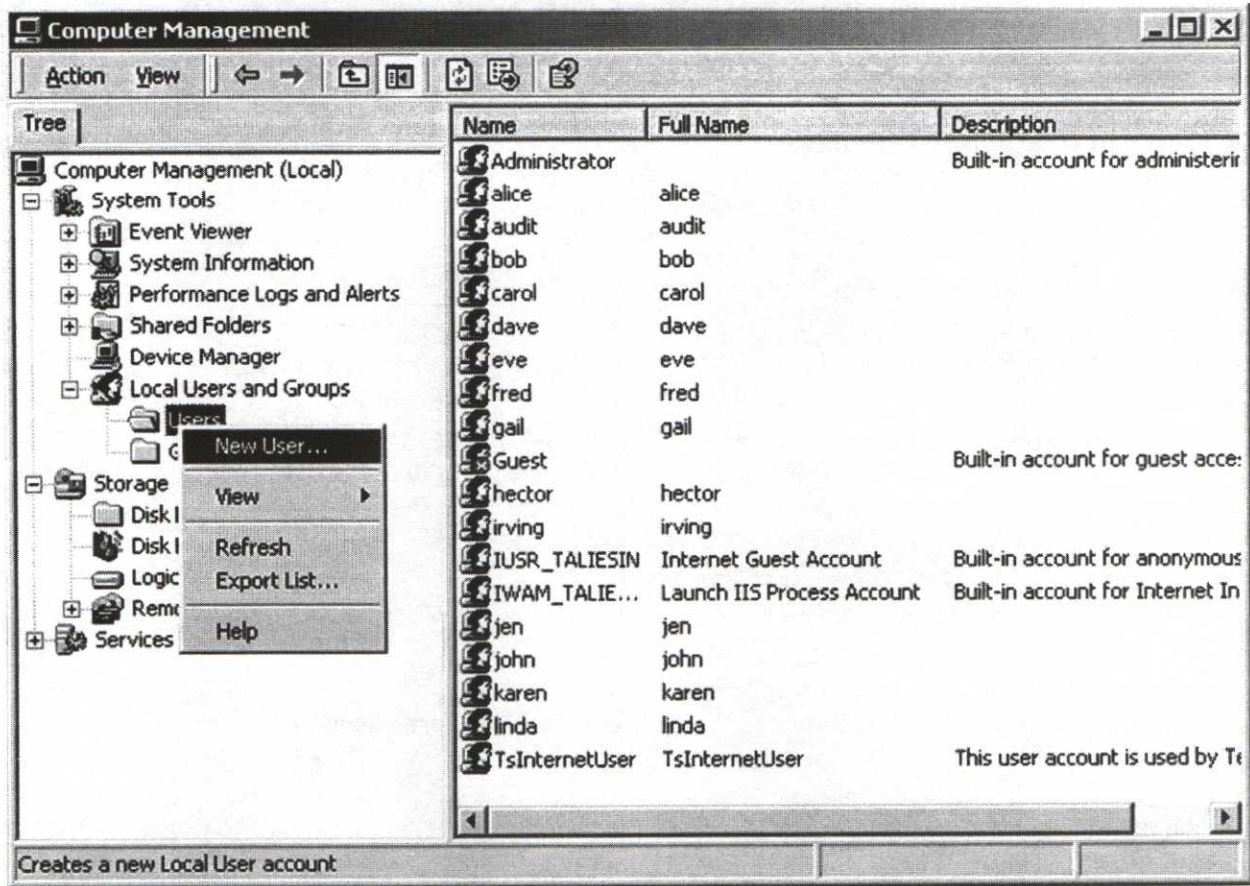
Part 1 – Install Cain and Abel

1. Locate the Cain and Abel setup file in the “Cain and Abel” folder in the “Tools\Windows Tools>Password Testing” directory.
2. Run ca_setup.exe to install the program. Accept the default installation options.

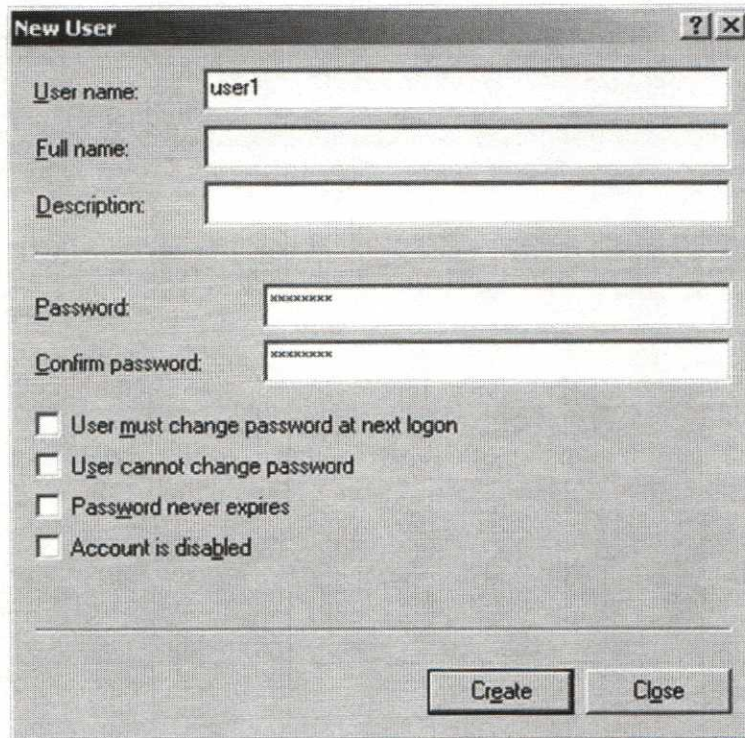
Part 2 – Create Sample Accounts on Windows

1. Open the Computer Management Microsoft Management Console (MMC). (Start → Programs → Administrative Tools → Computer Management.)
2. Expand the “Local Users and Groups” node.

3. Highlight the “Users” container. Right-click the “Users” container and select New User to create a new user account.



4. Create a user with a First Name and User Logon Name of **User1**. Enter the password “password” (without the quotes) for User1. Clear any/all of the check boxes.



New User [?] [X]

User name: user1

Full name:

Description:

Password: xxxxxxxx

Confirm password: xxxxxxxx

User must change password at next logon

User cannot change password

Password never expires

Account is disabled

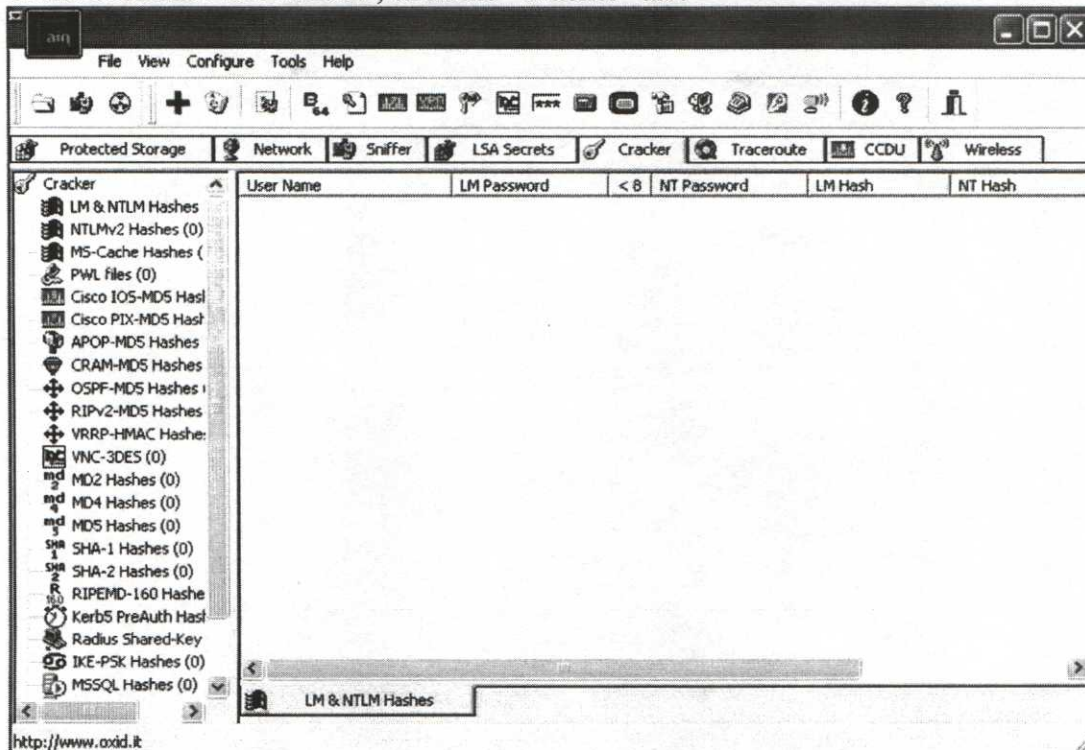
Create Close

5. Click **Create** to create the new account. (**Note:** If you receive an error message about the password when you try to create the account, you might need to modify your computer's password policy to allow the use of weak or blank passwords. Your password policy can be edited by going to Start → Programs → Administrative Tools → Local Security Policy → Account Policies → Password Policy.)
6. Repeat steps 3–5 to create at least four more user accounts (**User2**, **User3**, **User4**, and **User5**). Set passwords of varying strength for each user. For example, you might want to try:
 - Using a blank password.
 - Using a password of 7 characters or fewer.
 - Appending or prepending numbers to a dictionary password (i.e., password12).
 - Substituting numbers or symbols for common letters (i.e., pa55w0rd instead of password).
 - Using a password that meets your site's requirements for password complexity (for example, a minimum of 8 characters in length, including a mixture of uppercase letters, lowercase letters, numbers, and symbols).

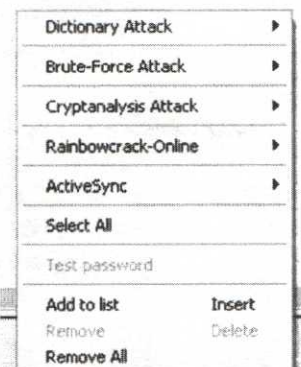
Part 3 – Use Cain and Abel to Extract the Local Password Hashes



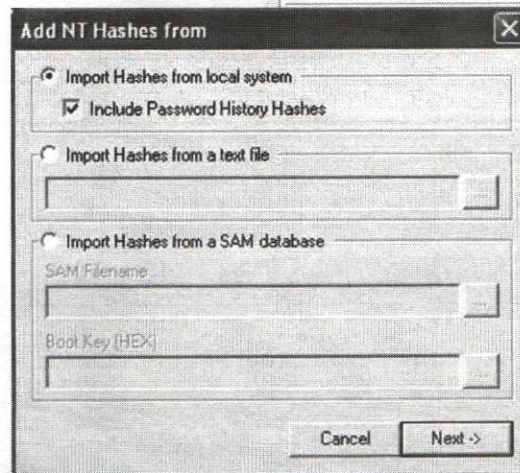
1. Start the Cain and Abel application (on the desktop).
2. In the Cain and Abel window, click the “Cracker” tab:



3. Select the “LM and NTLM Hashes” option in the left-hand pane.
4. Right-click in the body of the window and choose “Add to list” (see the figure to the right).
5. A dialog entitled, “Add NT Hashes from,” should now be visible on your screen. The first option, “Import Hashes from local system” should automatically be selected. Notice that you can also import from a text file or from a SAM database. A SAM database is typically the backup SAM database stored in the C:\Windows\Repair directory.



6. A second option allows us to attack the password history hashes as well. Turn this option on as well. Knowing the last password used can sometimes give us a really good sense for what the next password might be.



7. Right-click in the main body of the window and choose the “Select All” option. If you fail to do this, Cain will not allow you to crack anything because it won’t know what you want to do.
8. Before we can attack the passwords, we need a dictionary to attack them with. If you look in the “Tools\Wordlists” folder, there is an archive called “Dictionary.txt.” This is a sample wordlist file that we can use for some of our password assessments.
9. Once you have located this dictionary, please return to the Cain and Abel window. Right-click any of the selected users or passwords and choose “Select All.”
10. With the users all selected, right-click again and select “Dictionary Attack -> NTLM Hashes.”
11. The next dialog allows you to select a password dictionary to use. To add in a dictionary, right-click anywhere in the top pane and choose the “Add” option. To do so, after clicking the “Add” button, browse to your desktop and select the “Dictionary.txt” file that you found in the “Tools\Wordlists” folder.
12. After you have selected the word list, click the “Start” button. You should immediately see Cain go to work.

In the test case with the default options, Cain is able to break 2 of the 11 password hashes that we have in well under a minute. With better wordlists, case permutations, and so on, it will do even better.

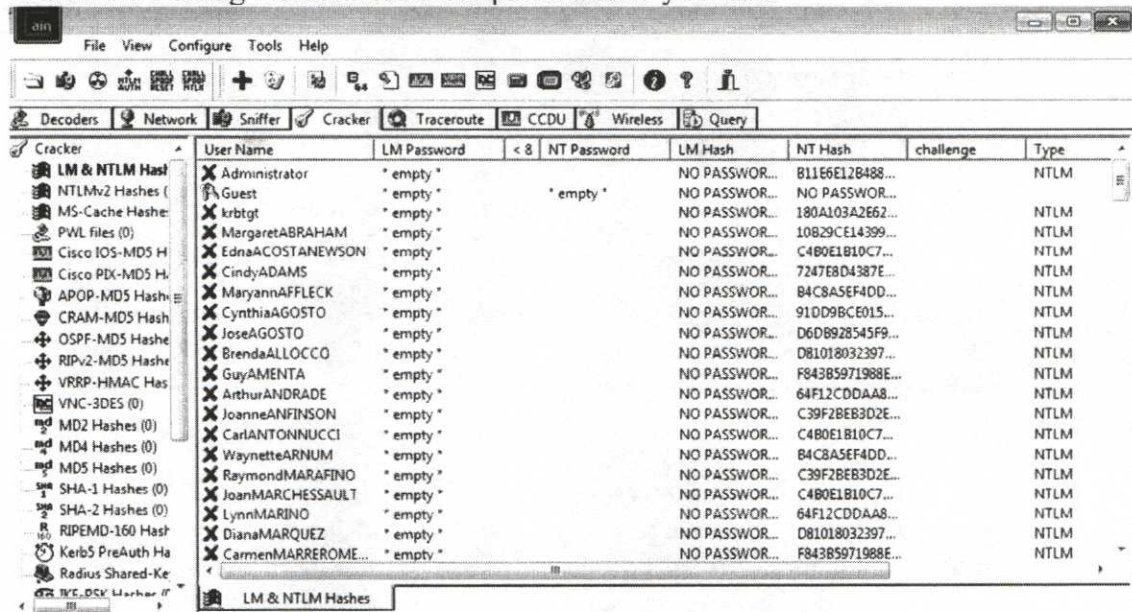
A very interesting result to notice is that when Cain and Abel starts, the brute-force attacks immediately after the dictionary attack, Cain very quickly determines the remaining pieces of three of the 11 passwords. The reason for this is that the remaining bits of the password are quite short, only being three or four characters long.

Part 4 – Domain Credentials and Cain

Let’s look at this again, this time with credentials off the Active Directory server in the lab domain. We’re also going to look at adjusting some of the settings in Cain to run a bit faster.

1. Look in the “Password Testing” folder in the “Windows Tools” directory. You should be able to locate a file named “507dc.pwdump.txt.” Prior to our arrival as auditors, we asked the administrators to provide us with the extracted password hashes as discussed in the course material.
2. Return to Cain. If you have just completed the last part of this lab, then you have all of your local accounts currently visible. Right-click and choose “Remove All.”
3. Now that the window is clear, right-click again and choose “Add to List.”
4. Rather than import accounts from our system, this time we will import from a file. Select this option, and then locate the “507dc.pwdump.txt” file from the

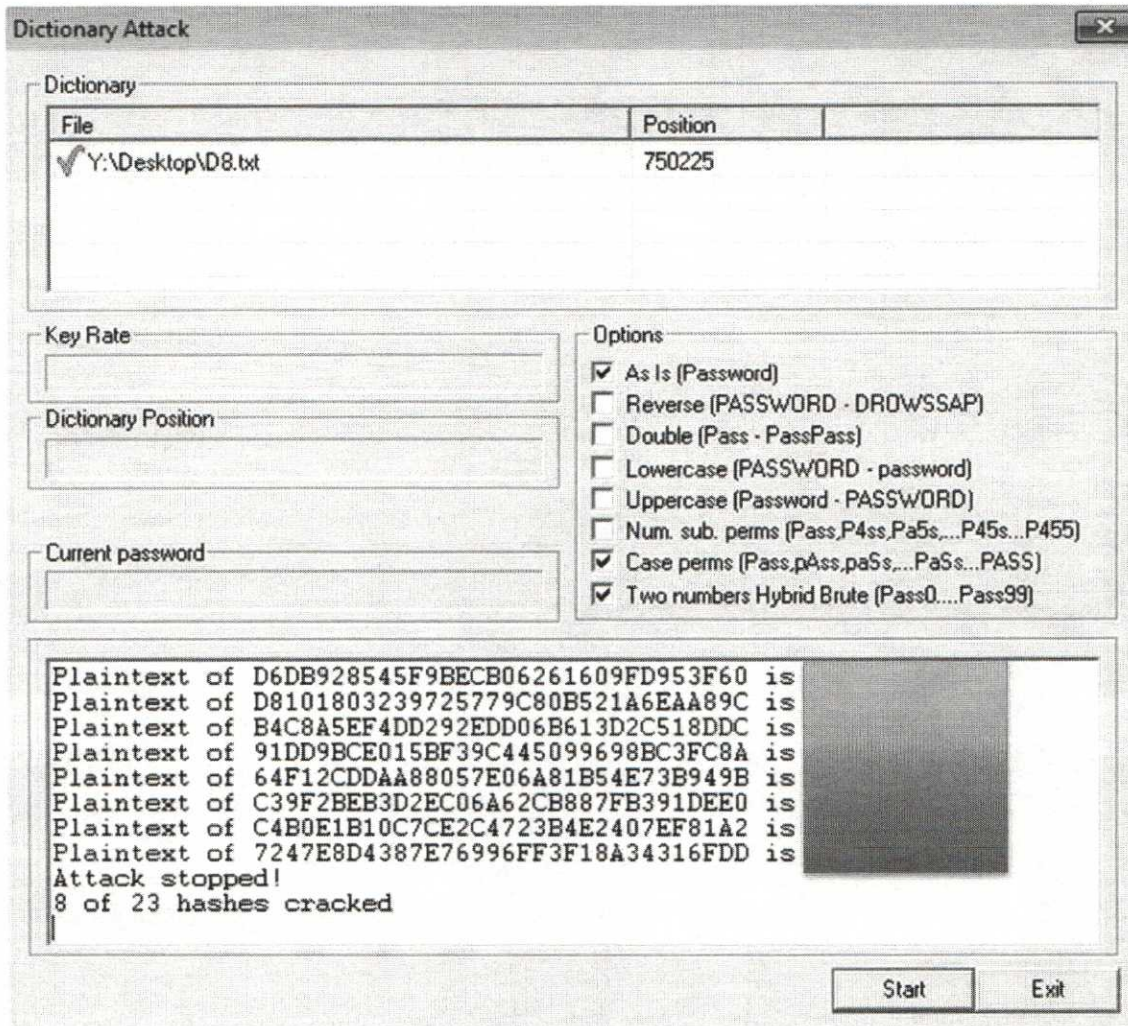
“Password Testing” folder. Add these passwords to your list.



<http://www.oxid.it>

- Now that these have been added, right-click and choose “Select All.”
- After all of the hashes are selected, right-click and choose “Dictionary Attack -> NTLM Hashes.”
- In the interface that comes up, Cain should “remember” your previous settings, including the dictionary file. **You cannot just start the test now!!** Cain not only remembers your settings, it also remembers which words from the dictionary have been tested so far! Right-click in the *dictionary* section and choose “Reset Initial File Position.” This will get us ready to run a fresh test.
- Next, adjust the password settings for the tests. We want to run it a bit faster. Please choose only the “As Is,” “Case Perms,” and “Two Numbers Hybrid Brute” options.
- Hit the “Start” button and wait!

After some time, you should start seeing results appear. If you are patient, you’ll see a number of passwords being cracked. The output that you see should look like the below, though we have concealed the cracked passwords:



There's one last thing that we'd like you to notice. Even though there were well over 100 users, there are only 23 hashes in use. What does this mean? This means that all of the users and computers put together have only 23 *different* passwords!

This isn't something that most people report on, but it can be very interesting. If you find 20,000 users but only 10,000 passwords, it tells you that many of the users have the same password as other users. Why would you care? Because it tells you that you're probably *not* selecting good passwords! You should find that passwords are unique with very few "collisions."

Protecting Data

Objective: Use tools to extract information about permissions and shares.

Estimated lab time: 20 minutes

Requirements:

- Laptop running Windows XP, Vista, 7, 8, or 10

Lab – Short Description:

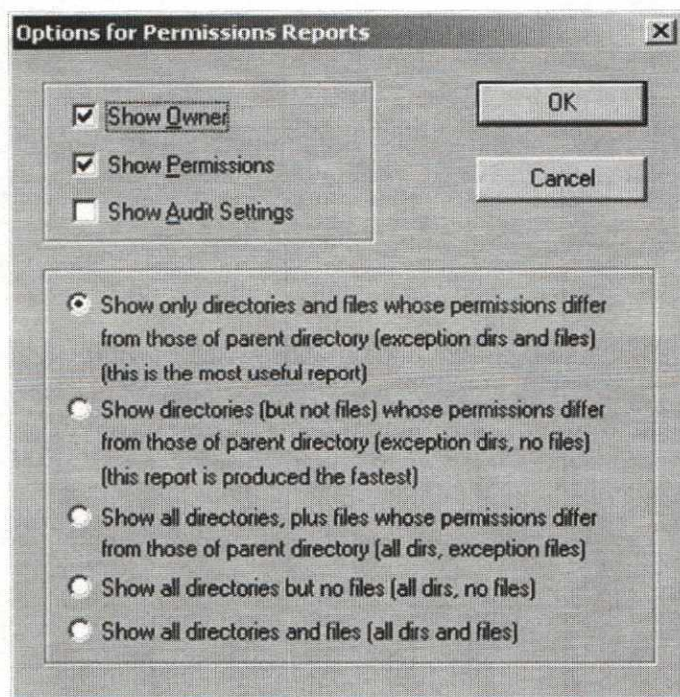
A short description of the lab steps is included for those students who are already familiar with Windows and related tools and are capable of working through the lab without detailed instructions. Those students who are less familiar with the subject matter should use the Detailed Description below, which includes step-by-step instructions.

1. Use Somarsoft's DumpSec to extract information about permissions and shares.

Lab – Long Description:

Part 1 – Use Somarsoft's DumpSec to Extract Information about Permissions and Shares

1. This lab makes use of DumpSec. The DumpSec installer can be found in the "Windows Tools" folder on your USB stick. After installing DumpSec, launch the DumpSec program.
2. From the **Report** menu, select the option for **Permissions report options**.
3. View the options available. Note that the default options show Owner and Permissions information. Also, the report will list the permissions for the root directory, and for any subdirectories and files whose permissions **differ** from the parent directory. In other words, DumpSec will assume that permissions are inherited, and show only objects where the permissions do not match the permissions on the parent object.

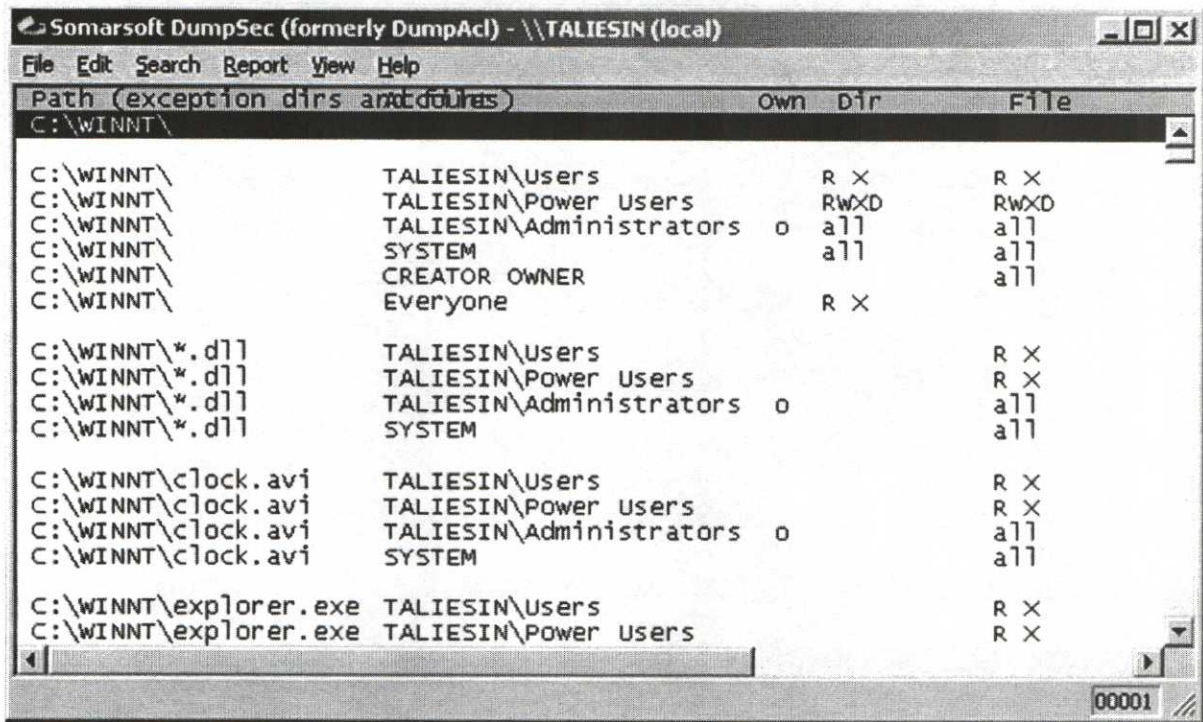


4. Click **OK** when you are done reviewing the options available.
5. From the **Report** menu, select the option for **Dump permissions for file system**. (Note: DumpSec can perform this same analysis against a remote system provided you are a local or domain administrator for that system. Simply select **Report, Select Computer** and specify the IP address of the target machine before selecting the permissions report.)
6. Select the **C:\Windows** directory to examine the permissions and click **OK**.



7. Click **OK** to start the scan.

8. When the scan is complete, review the results. You should see something similar to the following:



9. Note the permissions listed. What information should you look for in reviewing permissions on the system?

The following list is not comprehensive, but represents some items of interest:

- **Permissions granted to the "Everyone" group.** Everyone literally includes everyone – including anonymous web users (i.e., IUSR / IWAM account on an IIS server) and unauthenticated (null session) users. "Everyone" should be used sparingly, if at all; it is preferable to replace "Everyone" with Users or Authenticated Users.
- **Permissions granted to non-administrative groups.** The \winnt directory contains the majority of the OS files. These files should be protected from modification by anyone except for Administrators and other trusted users or groups. In most cases, non-privileged groups should require only Read (RX) or Execute (X) permissions to these files.

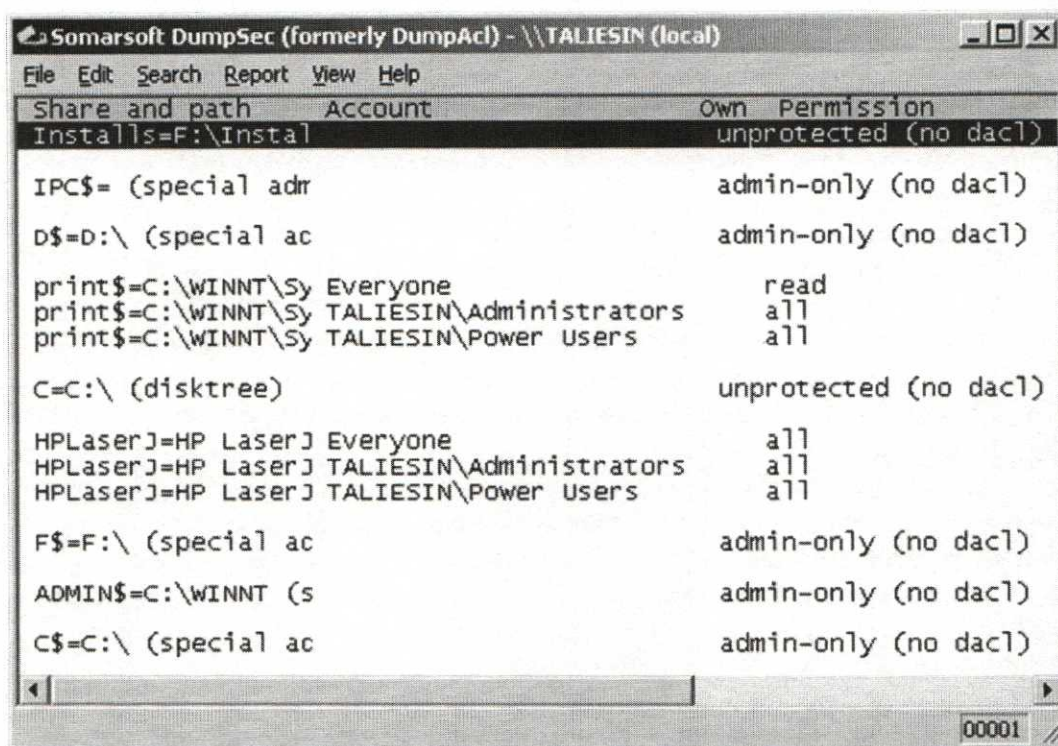
10. If you have a data directory or data drive on your system, scan that drive as well. What permissions are present?

Answers will vary depending on the directory scanned and the existing setup of the system. If the directory is a personal directory – i.e., My Documents – permissions will normally be limited to the current user only. If the directory is a data partition,

Windows will standardly set the permissions on the root of any drive to Everyone=Full Control.

If you are reviewing the permissions on a shared data directory – i.e., a file server, for example – you will most likely have an extensive set of permissions that will require the assistance of the data owners to review and assign appropriately.

11. From the **Report** menu, select the option for **Dump permissions for shares**. (Note: If you want to run this tool against the a remote target machine, select **Report, Select Computer** and specify the IP address of the target machine before selecting the share report.)
12. Review the results of the report. Note the permissions assigned to various shares; pay special attention to permissions granted to the “Everyone” group, or to shares marked as “Unprotected” (meaning the permissions are set to Everyone = Full Control. Note that such a setting would allow anonymous (null session) users to connect to the share and potentially access data within it, depending on the underlying NTFS permissions).



Security Configuration and Analysis

The following lab will be run **against your system** (the GUI tools used in this lab can be run only against the local system). Most of the following steps will simply gather information about your system; **however, the last part of the lab will involve applying configuration changes to your laptop.** If you have any concerns about applying the changes to your system, you can skip the last part of the lab.

Objectives:

- Learn to use Microsoft's Security Configuration and Analysis tools.
- Modify the security settings of a security template.
- Evaluate the security of an unknown system.

Estimated lab time: One hour

Requirements:

- Laptop running Windows XP, Vista, 7, 8, or 10

NOTE: If you are using Windows Vista or Windows 7, there are no default security templates that come with the operating system. Don't ask us why, because we don't know. ☺ If this is the situation that you are in, simply create a new template in the template MMC.

Lab – Short Description:

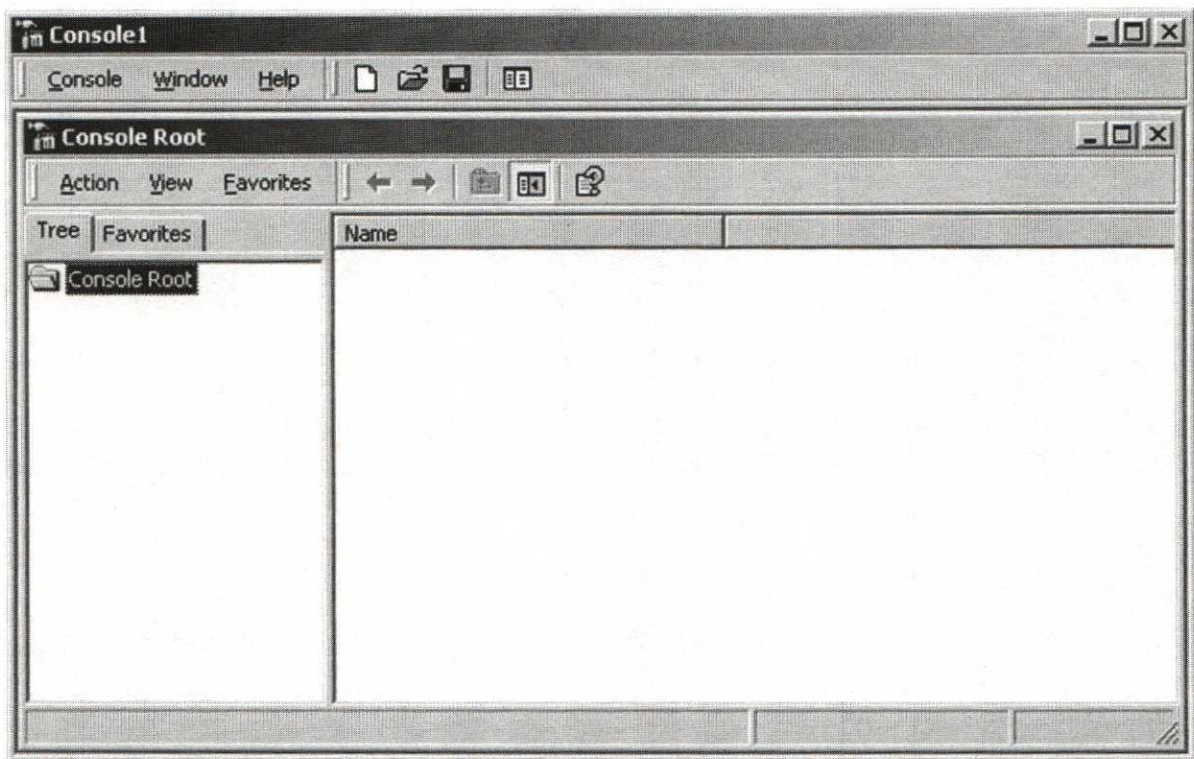
A short description of the lab steps is included for those students who are already familiar with Windows and related tools and are capable of working through the lab without detailed instructions. Those students who are less familiar with the subject matter should use the Detailed Description below, which includes step-by-step instructions.

1. Create a Microsoft Management Console (MMC) that includes the Security Configuration and Analysis and Security Templates snap-ins.
2. Expand the list of templates. Browse through a template to view the different security options available. Please note that Vista and Windows 7 do not have any security templates installed by default! If you are using Vista or Windows 7, simply create a new template as a starting point.
3. Modify one of the templates to reflect a specific security policy (the specific policy is listed in the Detailed Description section below).
4. Import the template into the analysis database.
5. Use the analysis database to analyze **your own system** (the GUI interface of the Security Configuration and analysis tool can be used only on the local host).
6. View the analysis results.
7. Customize the security template so it can be used to audit options not present in the default templates.

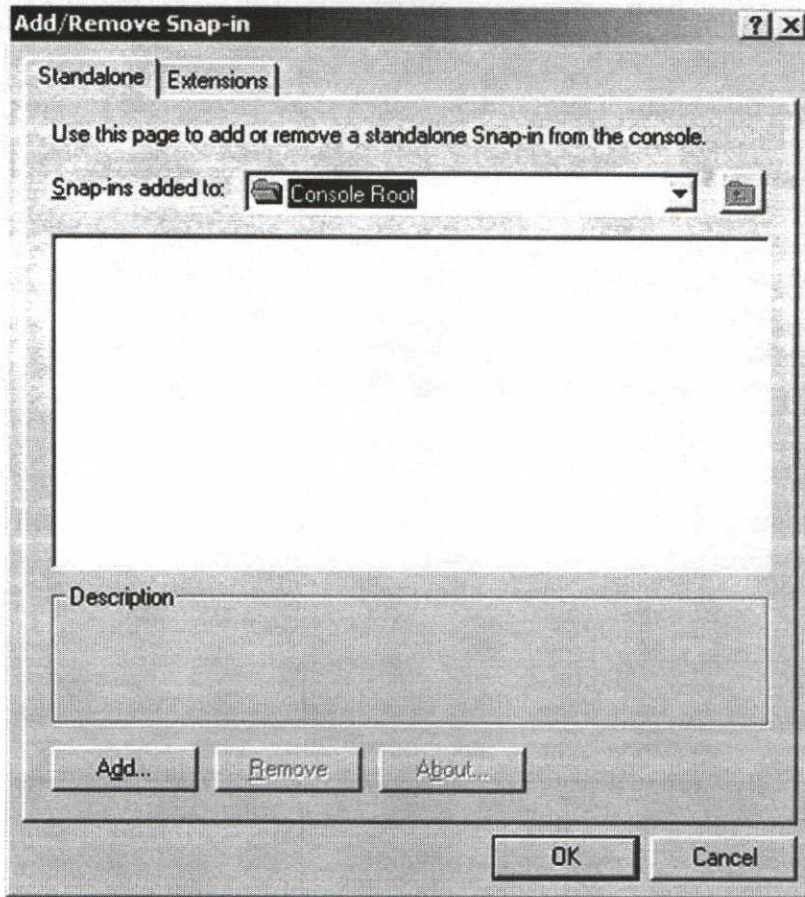
Lab – Long Description:**Part 1 – Create a Microsoft Management Console (MMC).**

*Windows includes a number of built-in MMCs with various default snap-ins. (You can find these in Windows by clicking **Start** → **Programs** → **Administrative Tools**. You might need to make the **Administrative Tools** folder visible by right-clicking the **Taskbar**, selecting **Properties**, selecting the **Task Bar** tab, clicking **Customize**, selecting the **Advanced** tab, and checking the box next to **Display Administrative Tools**.) You will create a new MMC that includes the **Security Configuration and Analysis** tools.*

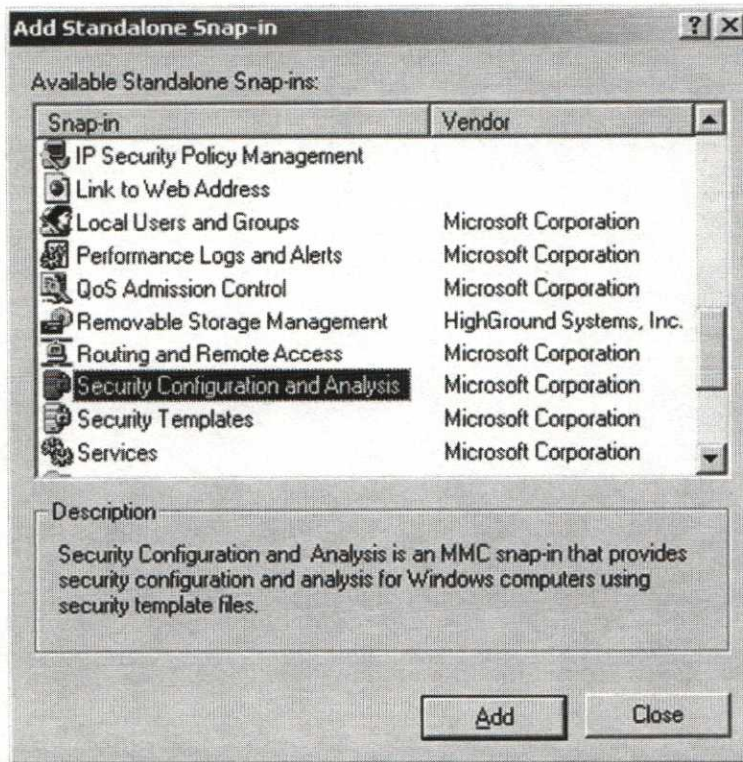
1. Click **Start** → **Run** → **mmc** and click **OK**. A blank Microsoft Management Console will appear.



2. Select **Console** → **Add/Remove Snap In**. The Add/Remove Snap-in Window will appear.



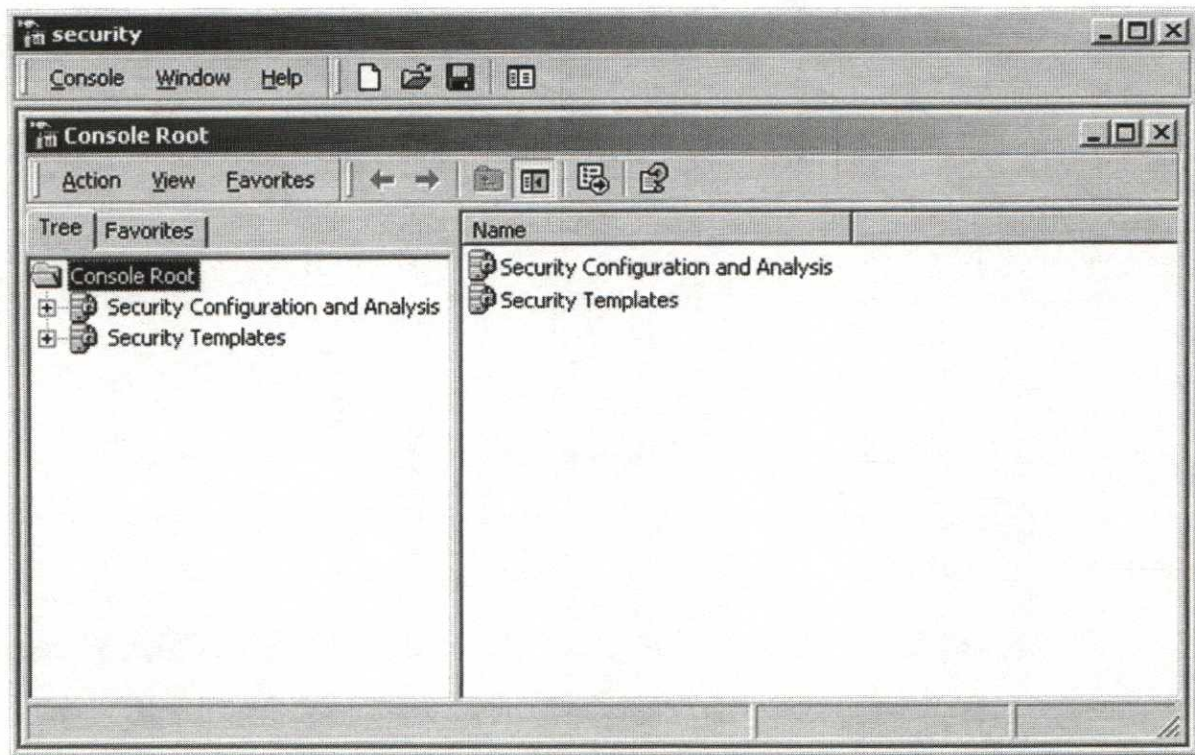
3. Click the **Add** button to display the Add Standalone Snap-in window.



4. Highlight **Security Configuration and Analysis** and click **Add**.
5. Repeat step 4 for **Security Templates**, and then click **Close** to close the Add Standalone Snap-in window. You should now see Add/Remove Snap-in window with your two selections added.



6. Click **OK** to close the Add/Remove Snap-in window.
7. You should now see your completed console. Select **Console** → **Save As** to save the file. Save the file as security.msc. By default, the console will be saved in the C:\Documents and Settings\- 8. You should now see your completed and saved console.

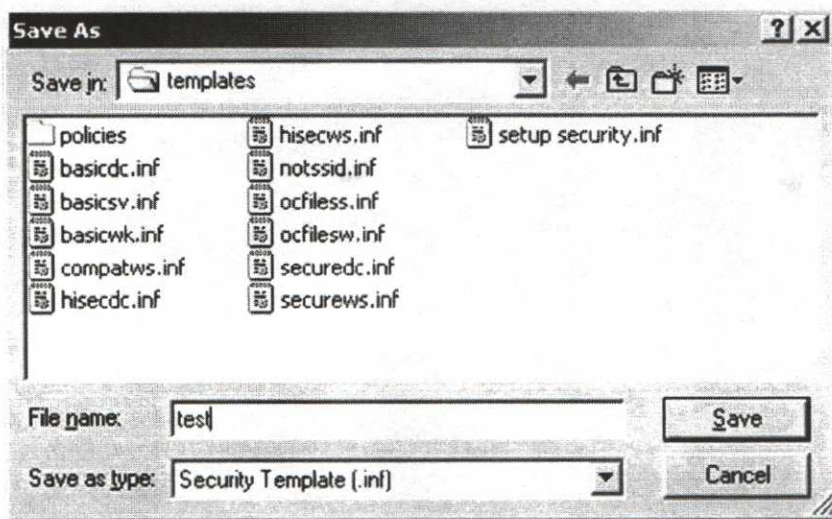


Part 2 – View the Available Templates and Options

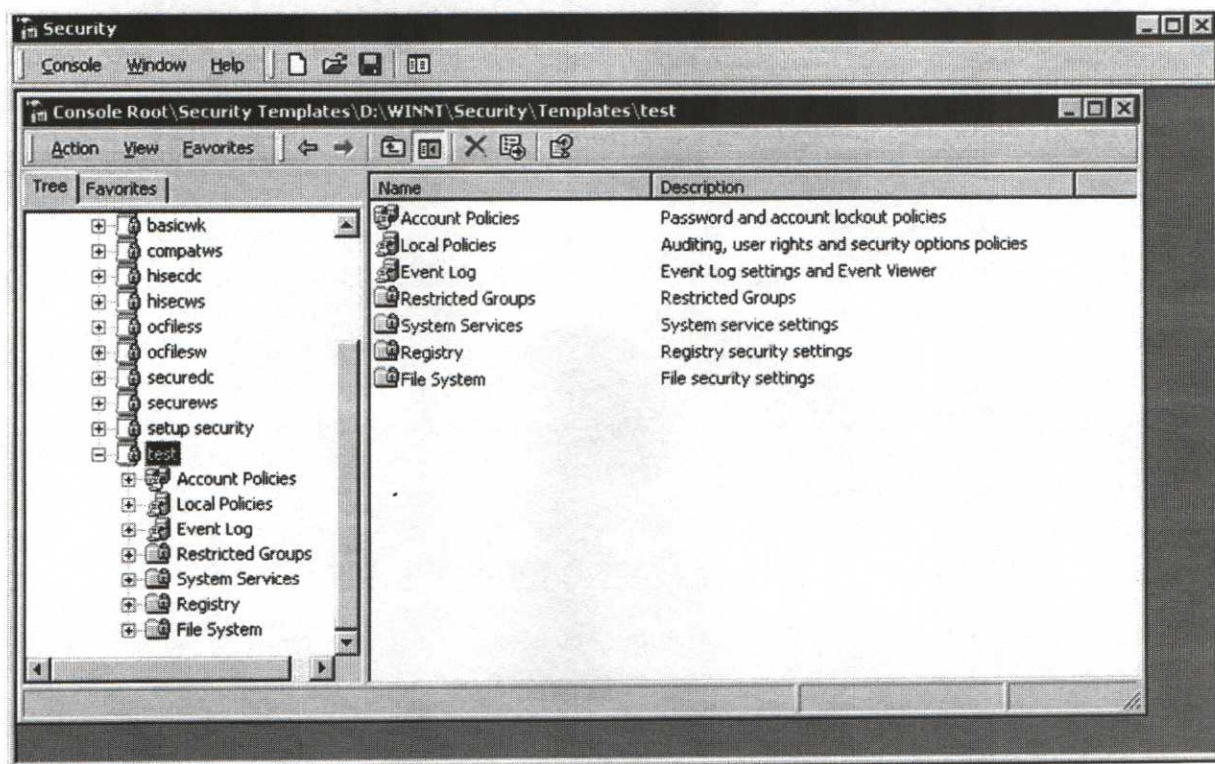
1. Expand the **Security Templates** node to view the available templates. These are the default templates included with Windows. Windows XP will show a number of available templates. Windows Vista and Windows 7 do not ship with any templates installed. If you are running one of these versions of Windows, please create a new template. Otherwise, expand one or more of the templates and browse through the various options that are available for configuration. Do not change any of the settings at this time.

Part 3 – Modify a Template to Match a Specific Security Policy

1. Either highlight one of the existing templates or create a new security template for your version of Windows.
2. If you are using an existing template, you might not want to modify the original template. Select the template that you'd like to use as a starting point, and then select **Action** → **Save As** to save the template. Name the template **test** and click **Save** to save the new template.



3. Expand your new **test** template. Your console should look like the picture below.

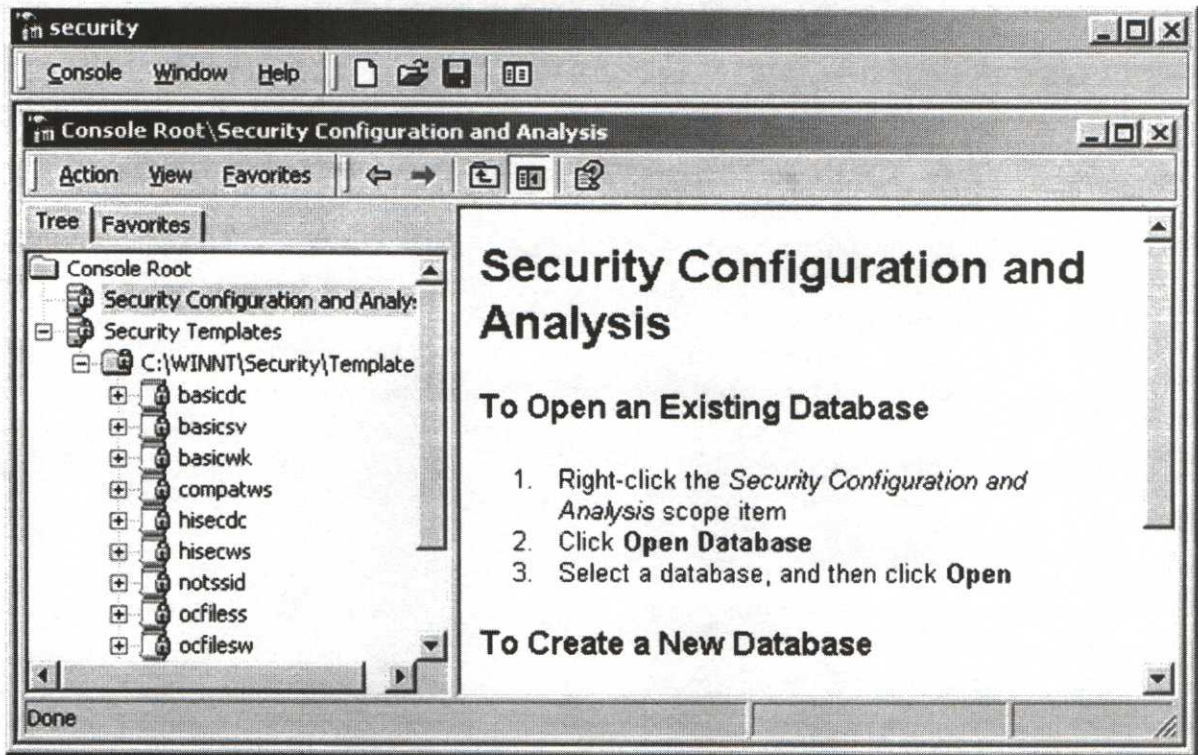


4. Modify the settings of your test template to reflect the following security policy:
- Account Policies → Password Policy
 - Passwords must be changed every 90 days; you must wait at least one day before you can change your password again.
 - Passwords must be a minimum of 8 characters and must meet complexity requirements.

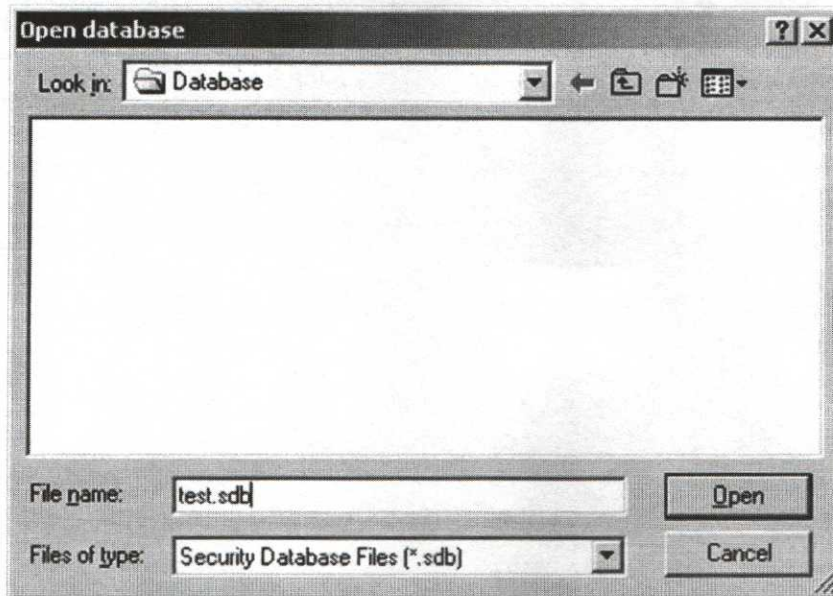
- System should remember the previous 10 passwords.
 - Account Policies → Account Lockout Policy
 - Accounts should be locked after 5 bad logon attempts, with the bad logon count reset after 30 minutes.
 - Accounts should remain locked for 30 minutes.
 - Local Policies → Audit Policy
 - Audit both successful and failed activities for the following:
 - Account management, logon events, account logon events, object access, policy change, and system events
 - Local Policies → Security Options
 - Under “Additional Restrictions for Anonymous Connections,” change the setting to “Do not allow enumeration of SAM accounts and shares.”
 - Enter a logon banner (message text) and logon title (message title) to warn users that the system is for authorized use only (you can make up your own text).
5. Once you have made the above changes, right-click the **test** template, and select **Save** to save the **test** template with your changes.

Part 4 – Import Your Template into the Analysis Database

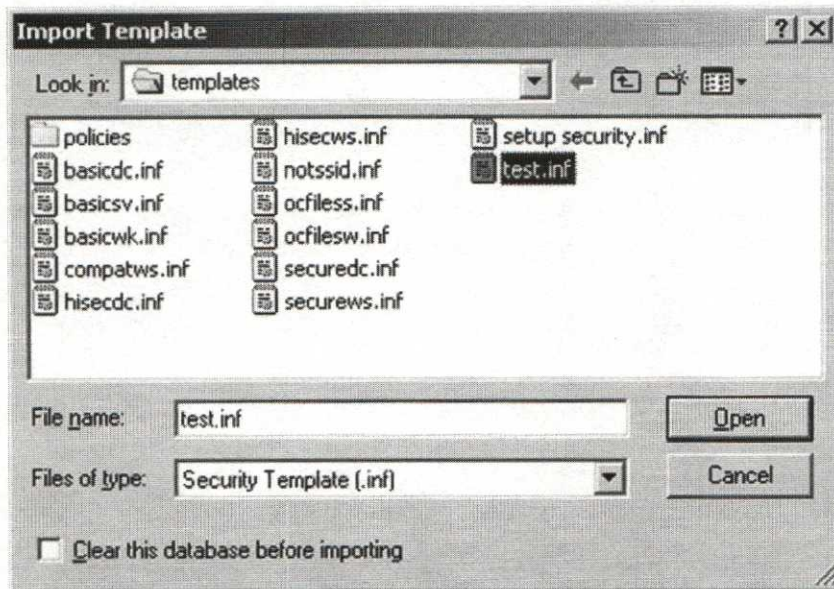
1. Highlight the **Security Configuration and Analysis** node. Your console should look like the picture below.



2. Right-click **Security Configuration and Analysis** and select **Open database**. You will need to select a name for the new database you are going to create. Enter **test.sdb** and click **Open**. By default, your database will be saved in C:\Documents and Settings\\My Documents\Security\Database.



3. You will be prompted to import a template into your new database. Select your **test.inf** template and click **Open** to import your template into the database.

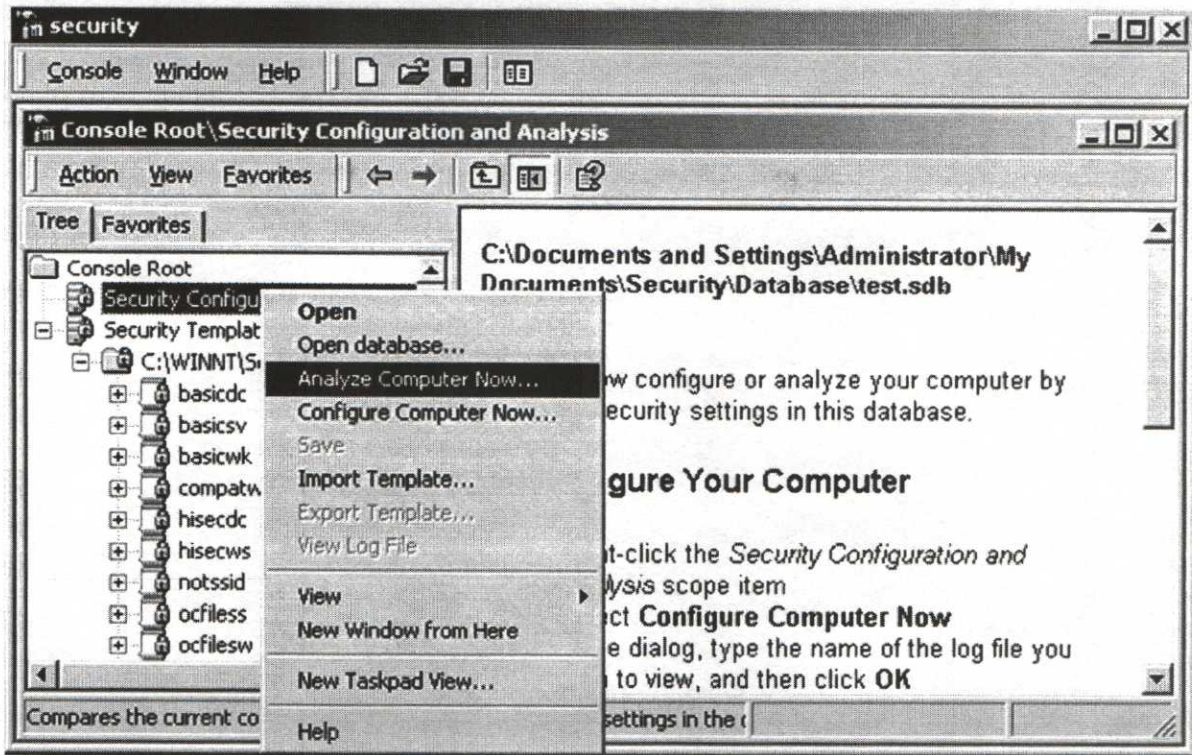


4. You are now ready to analyze your system.

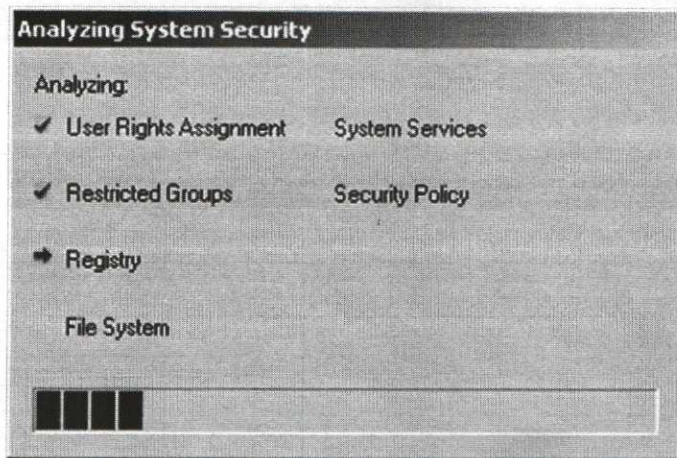
Part 5 – Analyze Your System

The following steps should be performed on your own local system. Running the analysis will simply read information from your system and compare it to the information in the database. Your system will not be modified.

1. Right-click **Security Configuration and Analysis** and select **Analyze Computer Now** to run the analysis on your system. Accept the default location for the log file (C:\Documents and Settings\\Local Settings\Temp\test.log) and click **OK** to begin.



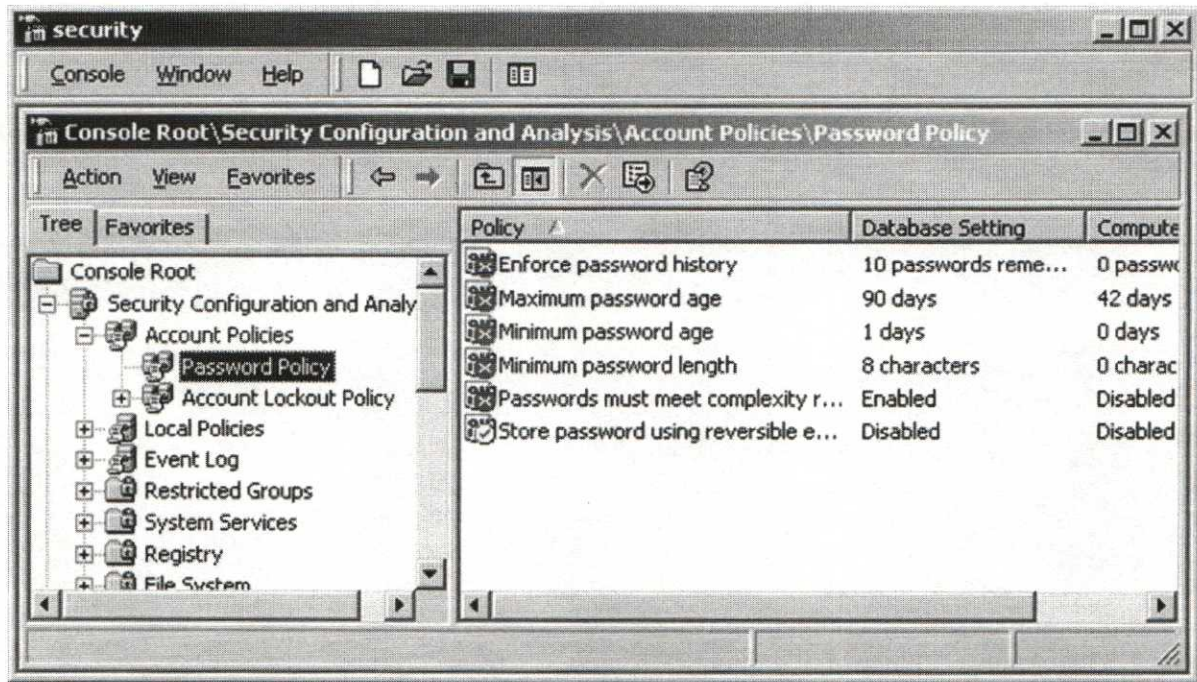
2. The analysis might take several minutes depending on your system resources. You will see a progress meter like the one shown below.



Part 6 – View the Analysis Results

1. After the analysis is finished, you should see a tree structure beneath the “Security Configuration and Analysis” node identical to the structure of your security template. Expand the various nodes to view the results of your analysis. A green check indicates that the system settings and the database settings match. A red X indicates that the settings do not match. If no icon is present, it means the setting is undefined

in your template/database and was ignored during the analysis. You should see something similar to the following.



2. Recall that a red X simply means that the settings do not match. It does not necessarily indicate that the existing system is less secure than the template/database settings.
 - Are there instances where a red X is present, but the existing settings are **more** secure than those in the database? (*Note: Answers will vary depending on the pre-existing configuration of students' laptops. One possible example of a mismatch would be maximum password age. The database we created specifies that passwords must be changed every 90 days. The default setting for Windows XP is to change passwords every 42 days. The default setting in this case could be thought of as **more** secure because passwords must be changed more frequently.*)
3. Browse through the remaining options to view the results.

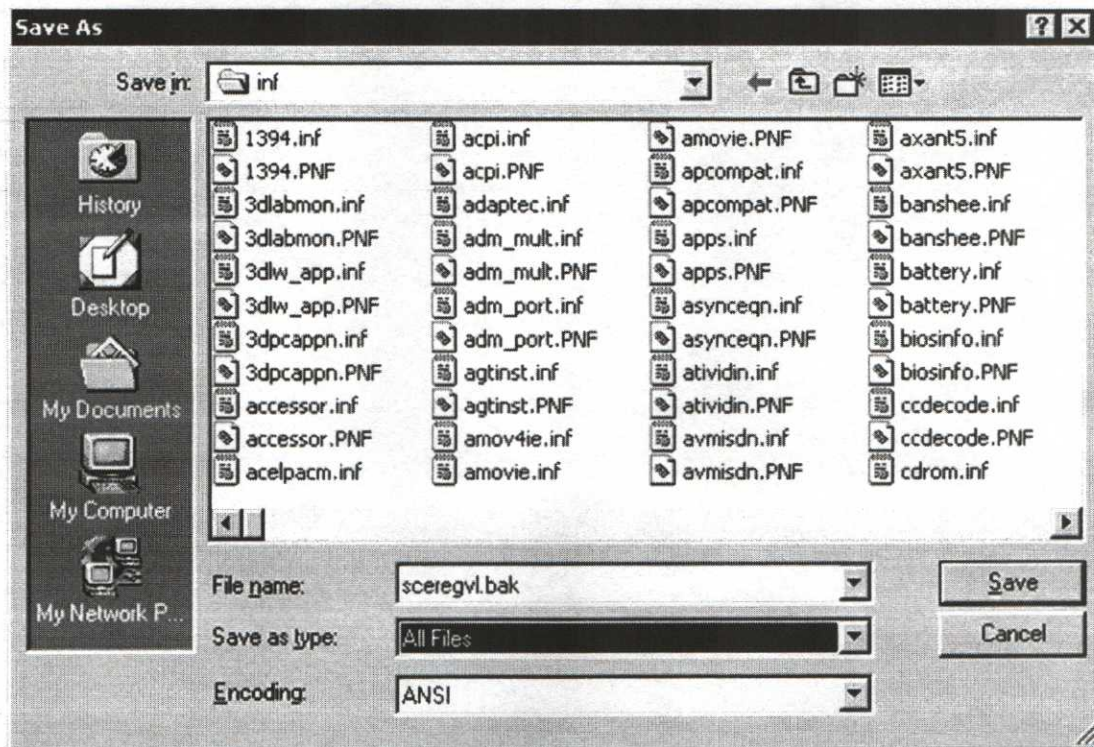
Part 7 – OPTIONAL: Customize a Template

If you have time remaining after completing the previous sections of this lab, you may continue with this section.

Note: The Microsoft Knowledge Base article 214752, “How to Add Custom Registry Settings to Security Configuration Editor”

(<http://support.microsoft.com/default.aspx?scid=kb;EN-US;214752>) provides detailed information on customizing security templates.

1. Locate the sceregvl.inf file in %systemroot%\inf. Double-click the file to open it in Notepad.
2. Create a backup copy of the file by selecting **File, Save As**. Name the file sceregvl.bak and click **Save**.



3. Close the sceregvl.bak file.
4. Re-open the sceregvl.inf file with Notepad.
5. You will add registry settings to allow you to disable the autorun feature on DVD-ROMs, and to disable storage of the LM password hash. When making the changes, **be sure to enter the information exactly as it appears below.**
 - **To add the autorun key:**
 - Locate the section of the configuration file that lists the entry for MACHINE\System\CurrentControlSet\Services\LanManServer...
 - **Above** that entry, add the following line:

```
MACHINE\System\CurrentControlSet\Services\CDrom\Autorun,4,%Autorun%,0
```

- The above line indicates the registry value to be set; the value type (4=REG_DWORD); the variable name for the text that will appear in the GUI (the actual text will be configured in the next step); and the display type (0=Boolean, or Enable/Disable).
- Locate the section of the configuration file that lists [Strings].
- **Below** that entry, add the following line:

Autorun = Disable Autorun on CD-ROM drives

- **To add the NoLMHash key:**

- Locate the section of the configuration file that lists the entry for MACHINE\System\CurrentControlSet\Control\Lsa\AuditBaseObjects...
- **Above** that entry, add the following line:

MACHINE\System\CurrentControlSet\Control\Lsa\NoLMHash,4,%NoLMHash%,0

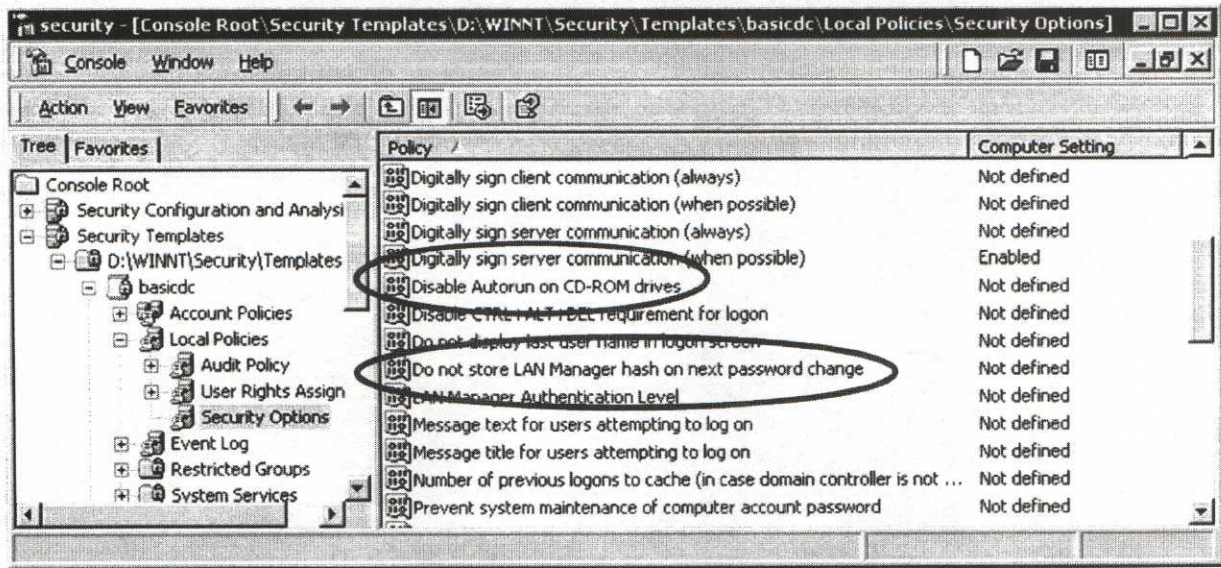
- The above line indicates the registry value to be set; the value type (4=REG_DWORD); the variable name for the text that will appear in the GUI (the actual text will be configured in the next step); and the display type (0=Boolean, or Enable/Disable).
- Locate the section of the configuration file that lists [Strings].
- **Below** that entry, add the following line:

NoLMHash = Do not store LAN Manager hash on next password change

6. Save the modified scereglv.inf file.
7. Register the changes you just made by running the following command at a command prompt:

regsvr32 scecli.dll
8. Exit the command prompt window.
9. Open your security.msc MMC (Start → Programs → Administrative Tools → security.msc).
10. Expand any template under **Security Templates**. Highlight Local Policies → Security Options.
11. Look for the two new entries you created (Disable Autorun on CD-ROM drives and Do not store LAN Manager hash on next password change). If you have done the lab correctly, you should see something similar to.

SANS Advanced Systems Audit Workbook



Day 5



All of the exercises in this section will be performed using “Unix Web Server” VMware image that was distributed by the instructor earlier in the week. Please start this system up now in VMware Player.

Many of the audit activities in today’s labs require that you act as the system administrator, running commands as root. To obtain a root shell, please log in using the information provided on the screen after the system boots. After logging in, you can become root by typing:

```
sudo su
```

You will be prompted for the password for the auditor account. Enter this password and your prompt should now be #, indicating that you are now root. Please remember that this system is “virtual” and you can always go back to the original version on the USB stick, so feel free to experiment! If you ever get into trouble or something ceases to function, try to restart the machine. If all else fails, unpack a new copy of it.

Section 1A

Exercise 1: Exploring Unix

The exercises in Section 1A map to the appendix material within the Day 5 book. To perform these labs, you have a choice when it comes to interacting with the UNIX system. You can either use the console session directly or you can log in remotely using Secure Shell. In practice, administrators rarely walk over to an actual console and log on. Instead, UNIX systems are traditionally administered remotely. For this reason (and others), we would strongly recommend that you connect to the system using the Putty Secure Shell client that is provided on the USB stick. **Please locate the “Putty.exe” tool and run it.** You will find this tool in the “Secure Shell – Putty” folder in the “Tools” folder that was on the USB Stick in the “Days 1-5” folder.

When you execute Putty, you will see a window very much like the image to the right. **Please use the IP address from your Unix Web Server to connect. When the connection completes, you will be prompted for the user ID to log in as. Please use “audit.” The password for this user is “Password1.”**

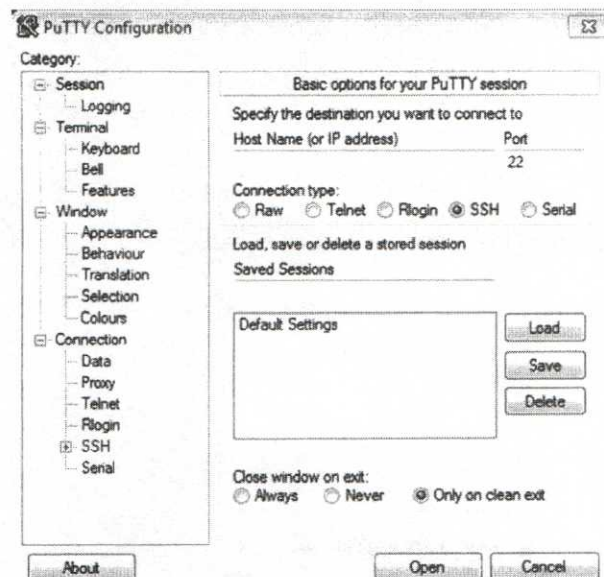
The Unix Manual: man

At this point, especially if you’re new to using Unix, we’d like to introduce you to the “man” or “Manual” command. Unix, from earliest times, has had the entire manual for all of the commands and other packages installed on the system in an online electronic format. These specially formatted files can be read in a number of ways, but the “man” command provides the easiest interface with some search capabilities.

To use this utility, you simply run the “man” command and give it another topic or command as an argument. The utility will then look up the manual page, if one exists, and present it to you one page at a time.

Please pull up the manual page for “ls” now.
(Hint: Type “man ls”.)

The “ls” tool, as you can see, is used to obtain lists of files. This is analogous to the “dir” command under DOS or Windows. There are a wide variety of options available for use with this tool and we recommend that you take a few moments to examine the possible options presented in the manual page.



What is the command-line option to obtain a “long” or detailed directory listing? Once you find it, please run the “ls” command with this option.

(Hint: Look at the command-line options in the manual page.)

(Hint: “-l” is the option.)

(Hint: To run this command, type “ls -l”.)

As you can see from the results, there is a lot of information displayed by “ls”. Using the manual pages, we can figure out what each option does and how to interpret the output. We’ll come back to the “ls” command in a little while. Let’s get back to the “man” command now.

The “man” command also has a range of possible command-line options. One form of that would include a number as the first argument to the command. For example:

```
man 2 umount
```

This command will look up the manual page for the “umount” utility, but what does the “2” do?

Please see whether you can identify what the number “2” in the manual command above stands for.

(Hint: “man man” might be of some assistance...)

(Hint: The Unix manual is broken up into “chapters” or “sections”.)

Try accessing the manual page for the “umount” command found in section “8” of the manual. Is it different from what is found in Section 2?

If you need to find a Unix command but did not know the command name, what syntax could you use with “man” to perform a keyword search of manual pages?

(Hint: “man man” will help you find the answer.)

(Hint: Read about the “-k” option.)

(Hint: “man -k topic” is what you’re looking for.)

(Bonus: Check out “man apropos”.)

Please run the command “man -k umount”. Can you see the two pages from Sections 2 and 8 that we looked at previously?

What would the command syntax be to do a keyword search for manual pages referencing “permissions”?

You might want to peruse the manual page for “ls” on the system to help answer the next few questions.

What are the numbers in the second column of the directory listing when you execute an “ls -l” on a Unix system?

Which option allows me to see all of the files in a directory even if they are hidden?

What does the timestamp associated with a file in a long directory listing document?

uname

The uname command can be used to display various facts about the system or operating system and it stands for “Microprocessor name.” The “u” at the front of the word stands for the Greek letter “Mu.”

Take a few moments to read the highlights of the manual page to familiarize yourself with the type of information that is reported and the format that is used. Although this information is somewhat basic, you can make excellent use of it in the course of assembling a baseline of the system.

What are the results when you execute the “uname -a” command?

What do these results mean? Why are they valuable in a baseline?

Unix Directory Commands: pwd, cd, and ls

Use the “pwd” command on the system to find out where you are in the file system. The “pwd” command stands for “Present Working Directory.” Which directory are we in?

Use the “cd” command (Change Directory) to switch to the root of the file system. (Hint: “cd” on its own will not put you in the root directory.)

(Hint: “cd” on its own moves you to your home directory.)

(Hint: To get to the root directory, enter “cd /”.)



(Note that in Unix, directory names are always qualified with the "/" (forward slash) rather than the "\" that we might be used to from Windows.)

In the root directory, use the "ls -l" commands to "list" the names of the files in this directory. What differences do you notice?

(Hint: Who owns the files? What kinds of files are they? Don't forget that you can use the manual page for "ls" to identify any unfamiliar elements in the output.)

It might be more helpful for us to see all of the files and precisely which of the files are files and which are directories. The "ls -la" command will help us with this. Please use this command now and take note of the differences again.

Take a few moments and poke around the file system a bit more. At this point, there's not much that can be done by you to harm the system. Another directory of particular interest is the "/etc" directory, which is where most of the configuration files for the system reside. One thing that you will notice as you use more and more UNIX systems is that although they are all very similar, the files might follow slightly different naming conventions and they might be in different locations.

For an auditing baseline, we would likely be interested in, at a minimum, taking a comprehensive inventory of the files in the "/etc" directory and in the default binary or "bin" directories. Simply having a listing of files is useful, but verifying whether or not the files have changed would be even better. We will examine this in some detail with the file integrity-checking exercise.

How can you take a full listing of all files on a Unix system, tracking modification times, ownership, permissions, and sizes?

(Hint: 'man ls' can be used to find command line options)

(Hint: "-R" can be used on most Unix systems to recursively list directory entries.)

(Hint: It is possible to redirect Unix command-line output to a file using the ">" followed by a filename.)

Move to the /tmp directory and create a file using your first initial and last name as the filename that contains a complete listing of all files and directories in the “/etc” section with file dates and sizes.

What command line did you use to do this?

Use the Unix manual to look up the “chmod” command. What purpose does this tool serve?

How can you change the permissions on a file so that the owner would have read-and-write access?

(Hint: There are two ways to do this.)

(Hint: Remember that Octal stuff we talked about in the appendix of the course book?)

(Hint: You can also specify the change using mnemonics.)

(Hint: The mnemonic form is “u+rw”.)

(Hint: The octal form is “600”.)

What side effect comes along with specifying the permissions using the octal notation?

(Hint: Can you change just the owner permissions using octal?)

Two other extremely important tools are “chown” and “chgrp.” Please look up each tool in the Unix manual.

What is the purpose of the “chown” tool?

What is the purpose of the “chgrp” tool?

How can either of these tools be used to change the owner or group for all files in a directory tree?

(Hint: Look for a “recursive” option.)

Also of note are the tools “cat,” “tail,” and “more.” Please look up the manual page for each of these tools to determine their function.

What is the purpose of the “cat” tool?

*In addition to “gluing” or concatenating files together, what else can it be used for?
(Hint: Try running “cat /etc/passwd”.)*

What is the purpose of the “more” tool?

Obviously, although “cat” is useful, quite frequently files will be too long to view using “cat”! The “more” tool was created for exactly this reason.

What is the purpose of the “tail” tool?

How can the “tail” tool be used to follow the contents of a file?

What does this mean?

Exercise 1B: Basic Scripting

As we learned during the Windows material, the ability to create basic scripts is invaluable for a system administrator and for an auditor. As an administrator, it allows for automation of repeated tasks and continuous monitoring and alerting.

For an auditor, scripting is invaluable because it allows us to create repeatable tests in addition to modeling continuous monitoring solutions that can be demonstrated for administrators. As an auditor, especially within an enterprise that is generating vast amounts of data, scripting can be an invaluable technique for parsing that data down to meaningful reports.

To get started scripting within a UNIX environment, it is a very good idea for you to get a little bit of hands-on experience with an editor that you can be virtually guaranteed will be installed on every UNIX system that you interact with. There are two such editors.

The first editor, which we will not actually teach you to use because of the level of discomfort that it would create, is “ed.” “ed” is the original editor on the UNIX system and was designed to be used with teletypes. Because teletypes print one line at a time and then advance the paper, there is no notion of using the cursor to move up or back or any other type of “movement” that is typical of editors that you use regularly.

Why mention “ed” at all? Because the View editor, “vi,” is largely a visual version of “ed.” Over the next two pages is a quick reference that can be used to learn a little bit about it.

vi Quick Reference

ENTERING vi

vi name start vi editor with file name .
The file is created if it doesn't exist.

LEAVING vi

ZZ exit from vi, saving changes.
:q! exit from vi, discarding changes.

CURSOR POSITIONING

h moves left one character position.
j moves down one line.
k moves up one line.
l moves right one character position.
0 (zero) moves to the beginning of a line.
w moves right one word.
b moves left one word.
CTRL-u moves up 1/2 screen full.
CTRL-d moves down 1/2 screen full.
G moves to the bottom of the file.
nG moves to line number n .
CTRL-l clear screen and re-draw.

TEXT MODIFICATION

itextESC inserts text to the left of the cursor.
Insert doesn't cause the cursor to move;
text appears as it is typed, terminate with
ESC.
atextESC appends (inserts) text to the right of
the cursor, terminate with ESC.
RtextESC Replaces (overprints) characters at the
cursor position, terminate with ESC.
dd deletes the line the cursor is on.
ndd deletes n lines from the cursor position.

D deletes characters from the cursor position
to the end of the line.
x deletes the character at the cursor.
nx deletes n characters to the right of the
cursor.
u undo the last change.

PATTERN SEARCHING

/pat/ positions the cursor at the next
occurrence of the string pattern.

NOTES:

ESC represents the ESC key. Press the ESC key when
it is called for in the above commands.

SANS Advanced Systems Audit Workbook

CTRL- represents the CTRL key. Hold the CTRL key and press the following key simultaneously.

CURSOR POSITIONING

```
}      move down one paragraph.  
{      move up one paragraph.  
mx     save the current cursor position and label it  
       with the letter x. (x is any letter)  
'x    return to the cursor position labeled x.
```

TEXT MODIFICATION

```
dw     delete the next word.  
.     (period) repeat last change.  
A     append at the end of the current line.  
P     put back deleted line(s). Text deleted with D  
and dd commands may be pasted back with the P  
command. Text is pasted in before the cursor  
position.  
:a,bs/old/new/  
      From line number 'a' to line number 'b',  
      substitute the pattern 'old' with the pattern  
      'new'. You may use any text string which  
      doesn't contain a carriage return in place of  
      the 'old' and 'new' strings. Use CTRL-G to tell  
      what line the cursor is on.
```

PATTERN SEARCHING

```
//     search for the next occurrence of a previously  
       specified search string.
```

MISCELLANEOUS

```
:w     write out current changes. The vi editor works  
       on a copy of your file. The :w command causes  
       the editor to write its copy over the original  
       which is on the disk.  
:w name write out changes to the file name . This is  
       like the :w command but the changes are written  
       into the file you specify. (good for making  
       intermediate copies)  
Cut and Paste Move to the beginning of the text to cut. Use  
dd to delete (cut) several lines. Use D to cut  
only the end of one line. Move to the place  
where you wish to paste the text. Use P to  
put back the text. You may need to clean up  
the spacing after pasting.
```

Additional help is available in **man vi**.

*(Vi Quick Reference chart referenced from
http://vertigo.hsrl.rutgers.edu/ug/vi_gref.html)*

Using the quick reference sheet and perhaps a brief tutorial led by your instructor, you should be able to use “vi” to create and edit a file.

Please use “vi” to open a file named “testscript.”

After “vi” opens, use the “i” command to switch to insert mode.

Add in the “shebang,” telling the shell to run the contents of this file through /bin/bash.




Figure 19 - Setting up the shebang

Next, let’s create a simple script to grab the default run level out of the system configuration files. To do so, we must use a few of the techniques that were covered in the slides.

VERY IMPORTANT: We are about to write a script that can find information from two different styles of UNIX systems. Of course, you have only *one* kind of UNIX system at your disposal right now. Why are we doing this? Because we want to create scripts that can be used on many systems rather than creating scripts that can be used in only one place. *This means that as you work through this lab, there will be files that you cannot find on the virtual machine!!!* This is completely expected. If you continue in the lab, you will get to files that you do have. Again, the script that we create will work not only on the system that you *have* but on other systems that are using the classic System V style startup scripts as well.

First, let’s take note of where the information will be located on various systems. Many systems will have a line that begins with the letters “id” in the file /etc/inittab:

```
linux:~# more /etc/inittab
# /etc/inittab: init(8) configuration.
# $Id: inittab,v 1.91 2002/01/25 13:35:21 miquels Exp $

# The default runlevel.
id:2:initdefault:

# Boot-time system configuration/initialization script.
# This is run first except when booting in emergency (-b) mode.
si::sysinit:/etc/init.d/rcS
```

Figure 20 - The “id” value in /etc/inittab

Other systems, particularly those using Upstart, will have the default run level configured in the `/etc/init/rc-sysinit.conf` file:

```
# rc-sysinit - System V initialisation compatibility
#
# This task runs the old System V-style system initialisation scripts,
# and enters the default runlevel when finished.

description      "System V initialisation compatibility"
author           "Scott James Remnant <scott@netsplit.com>"

start on (filesystem and static-network-up) or failsafe-boot
stop on runlevel

# Default runlevel, this may be overridden on the kernel command-line
# or by faking an old /etc/inittab entry
env DEFAULT_RUNLEVEL=2
```

Figure 21- DEFAULT_RUNLEVEL under Upstart

We would like the script to automatically identify these two files and extract the run level from them. Essentially, the test would be:

If the file `/etc/init/rc-sysinit.conf` exists then

Extract the number immediately following the equals sign on the line that begins “`env DEFAULT_RUNLEVEL=`”

Otherwise, if the file `/etc/inittab` exists then

Extract the number in the second colon separated column from the line that begins ‘`id`’

Print out the default run level at the console.

What we’ve just written above is known as “pseudocode.” All that we’ve done is written out, in English, what it is that we’d like to accomplish. Now we just need to translate it.

Please test yourself before turning the page! See how much you can figure out based on the material covered in the course book. Remember that you can test your code using the UNIX system! When you have either a working solution or have given up, turn the page.

```
#!/bin/bash
if [ -e /etc/init/rc-sysinit.conf ] ; then
    RUNLEVEL=`cat /etc/init/rc-sysinit.conf | awk -F= '/env DEFAULT_RUNLEVEL
/ {print $2;}'`
fi
if [ -e /etc/inittab ] ; then
    RUNLEVEL=`cat /etc/inittab | awk -F: '/^id/ {print $2;}'`
fi
echo The default runlevel for this system is $RUNLEVEL
```

In the script above, you might see something that doesn't look familiar. Notice that on the lines where we are assigning a value to the variable RUNLEVEL, the command to be executed on the right-hand side is surrounded by back quotes (` characters or accent grave). This tells the shell to first perform whatever actions are in the back quotes (execute that command), and then place the output into the variable, providing the result that we want!

To execute your script, you have two options. The brute-force method is to type the following:

```
bash testscript
```

A better solution is to mark the script as executable. As you should be well aware, we can accomplish this in the following two ways:

```
chmod 755 testscript
```

Or

```
chmod oug+x testscript
```

Now we can simply type `./testscript` to execute it!

Section 2

Exercise 1: File Integrity Checking

Purpose: Allow you to see a file checker in action.

As we've discussed, file integrity checkers are an excellent form of host-based security as far as intrusion detection goes. They are also quite powerful for figuring out what went wrong or which files have been modified if a security event has been detected.

As a reminder, when conducting an audit or trying to determine whether something has gone wrong, it is extremely important to have a copy of the baseline fingerprints for the files on the system stored on removable media like a CD-ROM, along with a static binary copy of the integrity checker to ensure that we are seeing accurate results.

When running this exercise, please be patient! File integrity tools tend to consume a lot of resources on the system being checked!

To begin, please switch to the "Unix Web Server" system in VMware and obtain a root shell prompt.

To use the file integrity assessment tool, we will need to act as the superuser on the system. Normally, we would need the administrator to either log in as root, use the "su" utility to become root or use the "sudo" utility to run the file integrity checker as root. The system has some file integrity software installed on it that we will experiment with.

The command to get us started is "tripwire." Run this command now, and then follow the directions to obtain more assistance.

What command-line option would be used to check the fingerprints of the files against the database?

(Hint: "tripwire -help" gives you more information.)

(Hint: "tripwire -m" is used to select the mode to operate in.)

(Hint: "tripwire -m c" runs a check against the current fingerprints database.)

When running Unix commands, it is usually possible to redirect the output through another utility. Try re-running "tripwire -m c" and pipe the output through "more." The command line will look like this:

```
tripwire -m c | more
```

Another excellent solution is to run the tool and "redirect" the output to another file. You can do this by running tripwire like this:

```
tripwire -m c > trip.report
```

It appears that there are a number of errors being generated. What do these errors mean?

(Hint: Tripwire says that it can't find certain files.)

Were these files modified or deleted since the fingerprint database was created? How do you know?

(Hint: If they were deleted or modified, they would show up in the report as "changed" or "missing.")

(Hint: How many files have changed?)

Use "CD" to change to the "/etc/tripwire" directory.

In this directory, you will find the file, "twpol.txt." This is the plain-text version of the policy for creating fingerprints. Looking through this file, can you see why Tripwire was reporting errors?

(Hint: The file and directories to fingerprint.)

(Hint: Are any of the files that generated errors listed in this policy file?)

(Hint: Because these files are explicitly defined, Tripwire will report their non-existence even though there is no fingerprint in the database.)

One of the handy things that you can do with this tool is automate it so that it will run periodically to verify that the system is remaining secure. This tool can quite easily be configured to run as a "cronjob." **Please take a few moments to read the manual page for "cron" and "crontab."** You will also notice that "crontab (5)" is referenced. This means is that there is another entry for crontab in Chapter 5 of the Unix manual. To read this man page, you would reference the chapter like this: "man 5 crontab".

What does the cron daemon do?

What do crontabs control?

How might you use a cron job to control Tripwire and for what purpose?

Exercise 2: Baseline Network Configuration

Purpose: Introduce additional tools that can be used to baseline the configuration and check the status of a system.

In this exercise, we will take a look at two more tools that can be used to examine the state of the system. In particular, we will take a brief look at the network configuration. The first tool that we will use is “netstat” or network status. *If you no longer have your VMware system running, please start it back up now and try out the “netstat” command on the system.*

As you can see, the “netstat” command gives you information about active connections as well as some statistics regarding those connections.

Please read through the “netstat” manual page briefly on the system and find out how to view the status of all sockets. What is the command line you would use to perform this check?

As with the Nmap tool we used previously, allowing the system to attempt to do hostname lookups might slow us down. What additional command-line option can we add to turn off this behavior?

Quite likely, the majority of the information that you requested scrolled off the screen. *Do you remember how we can pipe the output of this command through another command to slow it down? Please write down the complete command line to view the output of “netstat” one page at a time.*
(Hint: Look back at the discussion of “more.”)

What ports is the system listening on? What do these services do?
(Hint: Check /etc/services and try looking up what you find using a search engine.)
(Hint: The interesting ports are 587 and 25.)

As you can see, this tool is quite useful for determining what’s running on the system. We will look at another tool later that can be used to examine the process table, but the

process table does not always make it obvious that a program is listening on a port. “netstat” allows us to quickly baseline which ports are open from the machine level.

What would it mean if port scanner like Nmap tells you that a host is listening on a port but when you run “netstat” on that host, you do not see the port listening? (Hint: Under what circumstances would “netstat,” a standard Unix tool, fail to report that a service is listening on a port?)

(Hint: Who would have the power to replace the “netstat” tool with a different version that would only selectively report services?)

(Hint: When would an administrator choose to do this?)

(Hint: They wouldn't. An attacker who compromises the host will quite frequently replace binaries like “netstat” in an effort to conceal his or her presence!)

Let's take a look at a second tool, “lsof.” You might remember that the “ls” command is used to list files in a directory. The “lsof” command is used to list open files on the system. In a Unix operating system, everything, as far as is possible, is viewed and treated as a file. As a result, if we examine open files, we should also be able to see open ports on the system because they are represented somewhere by a file.

Let's begin by running the “lsof” command on the system.

“lsof” is not a default tool with many Unix operating systems; however, it is a freely available tool that you can easily install on any system or, more likely, put on your audit CD of statically compiled tools so that you can use it whenever you need it.

Please begin by skimming through the “lsof” manual page.

As you can see, it is quite long. We are interested in identifying network ports with “lsof,” so I'll share a trick with you.

If you enter “man lsof,” you will be at the top of the long manual page. To find what you are looking for, you can press the forward slash key (/), and then type a phrase. Please type “man lsof.” After the man page has been brought up, please use the “/” followed by “lstat”, and then press the Enter key. What happens?

Now, we're going to cheat because we know exactly what we're looking for. In this case, we'll search for "whose Internet address." This will drop us in the middle of the discussion of the option that we want to use. If you press the "b" key, it will take you back one page so that you can see what we're talking about... The "-i" option.

After you read through the description for "-i," try out "lsof" with the option. What happens?

What happens if you try "lsof -i tcp"?

You might not receive any output from this command. What could be wrong?

(Hint: Let's look at what's wrong. If "lsof" doesn't report back with information, chances are that you don't have enough access to find out which ports are connected to which processes on the system!)

(Hint: If you're logged in as root and still don't see any output, is it possible that there are no open ports?)

What information does this command tell you about each process?

Exercise 3: Startup Scripts

Purpose: Identify the location of the startup scripts on a Linux system.

Another area of the system that we'd like you to know about is in the "/etc" directory. We mentioned earlier that nearly all of the configuration for the Unix system will live somewhere in the /etc directory tree.

Right now, though, we'd like to look at the startup section. "/etc/rc.d" is where the startup scripts reside on most versions of Linux, but as we said, Unix is the same, but different. On the system that we are using in class, for instance, the "rc.d" directories are located directly off the "/etc" directory, and each of the files within those directories are links into the "/etc/init.d" directory.*

We are not going to dig through these files by hand because they are actually small programs. Instead, we want you to be aware of them and what they are for. These files control all of the aspects of how the various services running on the system will be started and, in some cases, whether they will be started at all.

A key file to know about is the "inittab" file on all System V derived Unix systems. This file is usually located in "/etc/inittab" and controls what run level the system will come up in. Run levels allow for various configurations of the system depending on what it's needed to do. (See "man init.") For System V derived systems, it will also typically define the behavior of some of the terminals connected to the system.

Please examine the "inittab" file in the "/etc" directory and determine the default run level for this system. What is the default run level?

(Hint: "more /etc/inittab")

(Hint: Don't try to hard... you will not be able to find what you are looking for!!)

Remember that not all UNIX systems are the same. Although the "inittab" is the typical place to find this information, the system that we are currently examining is being controlled by a mechanism known as "Upstart." In an Upstart driven system, this is controlled by the "/etc/init/rc-sysinit.conf" file. *Please examine the /etc/init/rc-sysinit.conf file and identify the default run level that this system will start at.*

(Hint: You are looking for the variable called "DEFAULT_RUNLEVEL.")

Now that we have the default run level, we can use this information to investigate which services will be started when the system reboots. Because the default run level is "2," the startup scripts that will be run are in "/etc/rc2.d." Looking at the names of these files, list at least three different types of services that will start when the system boots:

The next command we will look at briefly is the “ps” command. This command is used to list the process table, or all running programs. For this command, because you’ve been exploring the Unix systems for a while now, you will need to read the manual pages and experiment. Please attempt the following exercises, consulting the manual pages as necessary:

What will the “ps” command show if you give it no options?

What does the “ps -e” command do on the Linux system?

What does the “—forest” (2 dashes) option do on the Linux version?

Why would a snapshot of all running processes be valuable?

Exercise 4: Tiger

Purpose: To run and subsequently examine the output from the “Tiger” tool.

Tiger is an updated version of what was originally “COPS.” This tool will allow us to perform a host-based security assessment on the Unix system. We have already installed and configured (mostly) Tiger on the system that we are using for the labs. In practice, the installation and configuration of a tool such as this takes a significant level of expertise on the type of system being analyzed. For this reason, we strongly suggest that you work closely with the local Unix administrators in your organization if you choose to install or use this tool.

First, start up your VMware “Unix Web Server” system if you have not already done so. After doing so, please log in using the information provided on the console of that system.

Please take a few moments to read through the manual page for “tiger.”

Please run the “tiger” utility.

The next several questions are based on the information found in the report that Tiger generates. The report will be in “/var/log/tiger”. Quite likely, there will be a large number of reports in this directory.

Please use the “cd” command to switch to the /var/log/tiger directory. Using the “more” utility, review the log files.

What account permissions alerts does Tiger produce?

(Hint: Tiger will produce both Warnings and Alerts. For example, when examining the section where Tiger is checking accounts from /etc/passwd, it reports that Mail’s home directory has given Mail world write access. In this case, this really isn’t a problem because the mail program should own that directory.)

(Hint: As a suggestion, try to extract all of the records with the ALERT or FAIL markers to find the really important items.)

Even though ALERTs will be the most important, there is a very important issue that received only a WARN that relates to Root’s home directory. Can you see what it is?

If write permission has been granted for a group on a particular directory, what does this actually mean?

Reviewing the check of user accounts report section, what danger exists with regard to files in the “/root” directory?

Why is this potentially dangerous?

There are a number of files the Tiger finds that are not members of a “package.” Please read the manual page for “apt.” What does the term “package” mean in this context?

Why would it be important for an assessment tool to report files that do not belong to any packages?

Tiger creates a report listing all “set user-id” programs. What are “set user-id” programs?

Why would an inventory of these programs be an important part of any baseline?

Section 3

Exercise 1: Unix Log Files

Purpose: Examine various log files in the Unix environment.

Most of the log files that we will handle in this section will exist in the /var/log directory. You might want to go there now so you're ready.

The "utmp" file examines who is currently logged in. To examine this file, please try the "who" command on the system.

Examine the manual page for "who". What options provide reporting that might be useful for monitoring a system?

(Hint: The -idle option is useful for identifying accounts that have been "idle," especially accounts that have been idle for days!)

The "wtmp" file is where we keep track of who has been on the system. To examine this file, we use the "last" command.

Please read the manual page for "last". How could this program be useful to an auditor?

Why would it be valuable to track login times in addition to login lengths?

How could these records be useful for a forensic investigation?



Let's look at the system log as well. Now, the system log or "syslog" can get quite large, so we're not going to read through the whole thing. What we'd like to do is look at the last few entries in the file. The file, by the way, might be called "syslog" on lots of systems, but both of the systems we're examining now call it the "messages" file.

To examine the last few lines of this file, we can use the "tail" command. Take a quick glance at the man page and you'll see that it can be quite flexible.

Using the tail command, dump the last few records in the messages file. What does the fourth column in the log file represent? (Columns are separated by spaces for this exercise.)

Please remember that syslogs can be collected on a single system for consolidated analysis. How could this fourth column be useful for us?

Another file to take a quick look at is the dmesg file. The dmesg file stores system messages or kernel boot messages, depending on the system. To view the file, you will need to be root or use the “dmesg” command. Please read the manual page for dmesg now, and then try this command.

What sort of information does “dmesg” report?

Exercise 2: Log Analysis with SWATCH

Purpose: Introduce an effective tool for sifting through centralized log files for audit exceptions and a tool that can be automated for regular audit exception reports from centralized or distributed syslog files.

Introduction:

Now that we've spent some time perusing the Unix file system by hand and paging through system logs, we can see the need for an effective way to parse these logs and report on anomalies that appear in them. Of course, we will be able to use the same techniques discussed in this exercise to audit an actual message log for exceptions.

The tool that we will evaluate is called "Swatch," or "Simple Watcher." This tool is available from <http://www.oit.ucsb.edu/~eta/swatch/> as free software and is an outstanding retrospective and real-time analysis and reporting tool for log files, though the output is a bit raw. The output from this tool can easily be run through a reporting tool like Crystal Reports or something similar once exceptions are identified.

To begin, please start up your VMware "Unix Web Server" system if you are not already logged in.

Please use the "mkdir" command to create a temporary directory with the name of your choice in the /tmp. The directory that you create should have, as its first letter, your first initial. Please follow this with your last name. In other words, if your name is "John Doe," the directory that you create will be named "jdoe."

(Hint: "mkdir /tmp/jdoe")

After creating the directory, please use the "cd" command to switch to this directory. We will build the configuration files and report files here.

(Hint: "cd /tmp/jdoe")

Please use the "man" command to read the manual page for the "swatch" tool.

(Hint: "man swatch")

(Feel free to use the area below for notes on interesting features.)

As you can see, this tool is quite flexible and can be used to post-process a log file or watch a file in real time. Even more powerful is the ability to have Swatch run any command or set of commands when an event is detected.

To use Swatch effectively, we will need to write a configuration file. The configuration file is essentially a description of what we want to see or what we want not to see. Most auditors and administrators find it most effective to use Swatch to “ignore” known events so that all other events are alerted rather than trying to define every anomalous event.

The configuration file is essentially a list of keywords followed by options and/or values to define how Swatch should act or react. Let’s begin by looking at the first few lines in the /var/log/messages file and begin excluding normal stuff.

Use the “more” command to list the first screenful of lines from the /var/log/sample_messages file.

(Hint: “more /var/log/sample_messages”)

What do the messages in the first screenful of output have to do with? In other words, what event took place to create the messages in the file?

(Hint: When does the syslog server start?)

(Hint: The log will typically begin with the system rebooting or with a log rotation. In a log rotation, you will typically find the first entry listed as the syslog server starting.

Why might we choose to ignore the CRON daemon log entries?

Why might we choose NOT to ignore the syslog daemon restarting?

For our purposes today, we will not need to report on all of these. The system restart line is sufficient. To exclude them, we can use the “ignore” keyword with Swatch. Let’s begin by creating a file on the system called “swatchconfig”.

1. vi /tmp/swatchconfig
2. Press the “i” key. This key puts the editor into “insert” mode.
3. Enter the following lines:

```
ignore /* Linux CROND .* run-parts ./
```

```
watchfor /*/  
echo
```

4. After entering these lines, Press the Escape key followed by “:x”.
5. At the command prompt, run this command:
**swatch --config-file=/tmp/swatchconfig --examine=/var/log/sample_messages |
more**

We will discuss what we did in the configuration file after we answer a few questions...

*What difference can you see between the output from the swatch command above and the “more /var/log/sample_messages” command that we entered earlier?
(Hint: Do you still see the normal messages?)*

Now how about that editor? Here’s a quick reference from a Rutgers site:

Vi Quick Reference

ENTERING vi

vi name start vi editor with file name .
 The file is created if it doesn't exist.

LEAVING vi

ZZ exit from vi, saving changes.
:q! exit from vi, discarding changes.

CURSOR POSITIONING

h moves left one character position.
j moves down one line.
k moves up one line.
l moves right one character position.
0 (zero) moves to the beginning of a line.
w moves right one word.
b moves left one word.
CTRL-u moves up 1/2 screen full.
CTRL-d moves down 1/2 screen full.
G moves to the bottom of the file.
nG moves to line number n .
CTRL-l clear screen and re-draw.

TEXT MODIFICATION

itextESC inserts text to the left of the cursor.
 Insert doesn't cause the cursor to move;
 text appears as it is typed, terminate with
 ESC.
atextESC appends (inserts) text to the right of
 the cursor, terminate with ESC.
RtextESC Replaces (overprints) characters at the
 cursor position, terminate with ESC.
dd deletes the line the cursor is on.
n dd deletes n lines from the cursor position.

D deletes characters from the cursor position
 to the end of the line.
x deletes the character at the cursor.
nx deletes n characters to the right of the
 cursor.
u undo the last change.

PATTERN SEARCHING

/pat/ positions the cursor at the next
 occurrence of the string pattern.

NOTES:

ESC represents the ESC key. Press the ESC key when
 it is called for in the above commands.

SANS Advanced Systems Audit Workbook

CTRL- represents the CTRL key. Hold the CTRL key and press the following key simultaneously.

CURSOR POSITIONING

} move down one paragraph.
{ move up one paragraph.
mx save the current cursor position and label it with the letter x. (x is any letter)
'x return to the cursor position labeled x.

TEXT MODIFICATION

dw delete the next word.
. (period) repeat last change.
A append at the end of the current line.
P put back deleted line(s). Text deleted with D and dd commands may be pasted back with the P command. Text is pasted in before the cursor position.
:a,bs/old/new/
From line number 'a' to line number 'b', substitute the pattern 'old' with the pattern 'new'. You may use any text string which doesn't contain a carriage return in place of the 'old' and 'new' strings. Use CTRL-G to tell what line the cursor is on.

PATTERN SEARCHING

// search for the next occurrence of a previously specified search string.

MISCELLANEOUS

:w write out current changes. The vi editor works on a copy of your file. The :w command causes the editor to write its copy over the original which is on the disk.

:w name write out changes to the file name . This is like the :w command but the changes are written into the file you specify. (good for making intermediate copies)

Cut and Paste Move to the beginning of the text to cut. Use dd to delete (cut) several lines. Use D to cut only the end of one line. Move to the place where you wish to paste the text. Use P to put back the text. You may need to clean up the spacing after pasting.

Additional help is available in **man vi**.

(Vi Quick Reference chart referenced from http://vertigo.hsrl.rutgers.edu/ug/vi_qref.html)

We don't really need to know much about *vi* to use Swatch, but we need an editor of some kind. There is an odd syntax that we need to know about in order to use Swatch that is mentioned in the manual page and which we used in the beginning configuration file. The manual page makes reference to something called a **regex**.

Regex stands for "Regular Expression." A regular expression is a method for defining how some text will look in a general way rather than a specific way. In other words, we can define a pattern that looks for any three characters followed by a number between one and four that appears at the end of a line:

```
./*[a-zA-Z][a-zA-Z][a-zA-Z][1-4]$/
```

If we were to include this in one of the "watchfor" rules in Swatch because it's a known exception pattern, we could easily generate a very specific report. More usually, we will use regular expressions to try to define in a general way messages in which we have no interest. For instance, sendmail usually reports mail events (like delivering or receiving email) to the messages file. It is quite likely that we don't want to see these messages in the reports, but every one of them will be different. Rather than trying to define a specific email message to exclude, we define a general pattern of message to exclude, which is where regular expressions come in. Here are a few basics:

<u>Pattern</u>	<u>Meaning</u>
.	Matches any single character
[<i>set or range</i>]	Matches anything in the set of characters within the braces or in the range defined in the braces. Examples follow.
[a-z]	Matches any lowercase letter between "a" and "z."
[0-9]	Matches any number between "0" and "9."
[aeiou]	Matches any vowel contained in the list within the braces.
*	Matches any number of occurrences of the previous expression (including zero occurrences).
+	Matches any number of occurrences of the previous expression > 0.
^	Matches the start of the line.
\$	Matches the end of the line.
?	Matches any single character.

SANS Advanced Systems Audit Workbook

Let's try to define a few regular expressions. Please keep in mind that there are many ways to define the same expression, so don't be upset if your expression isn't identical to the answers on the next page. If you are unsure whether your answers are correct, please ask the instructor or a proctor for some assistance to check your work.

It is important to note that some applications can find a regular expression within a line without the need to define surrounding patterns, whereas other tools require the entire line to be defined. This means that it might be necessary to define an expression using `".*"` around your actual target expression in order to match the remainder of the text in the line.

Please define a regular expression to match each of the following:

1. At least one number followed by an uppercase letter between "N" and "R."
2. Modify the last expression so that it can be found in the midst of a larger string.
3. A line beginning with the letter "Z" containing a number between 10 and 99 and ending with the letter "Q". The rest of the line is unknown.
4. A letter between "A" and "Z" followed by a five-digit number.
5. Modify the last expression so that it can be found in the midst of a larger string.
6. Match the pattern "**Login failed:**" followed by at least three lowercase characters.
7. Modify the last expression so that it can be found in the midst of a larger string.

Please define a regular expression to match each of the following:

1. At least one number followed by an uppercase letter between "N" and "R."
`/[0-9]+[N-R]/`
2. Modify the last expression so that it can be found in the midst of a larger string.
`/.*[0-9]+[N-R].*/`
3. A line beginning with the letter "Z" containing a number between 10 and 99 and ending with the letter "Q". The rest of the line is unknown.
`/^Z.*[1-9][0-9].*Q$/`
4. A letter between "A" and "Z" followed by a five-digit number.
`/[A-Z][0-9][0-9][0-9][0-9][0-9]/`
5. Modify the last expression so that it can be found in the midst of a larger string.
`/.*[A-Z][0-9][0-9][0-9][0-9][0-9].*/`
6. Match the pattern "Login failed:" followed by at least three lowercase characters.
`/Login failed: [a-z][a-z][a-z]+/`
7. Modify the last expression so that it can be found in the midst of a larger string.
`/.*Login failed: [a-z][a-z][a-z]+.*/`

One additional point that can be rather important is that if you need to match any special character like a +, *, ., and so on, you must escape the character with a backslash. This is why the forward slashes used in the configuration file were preceded with a backslash.

Using the information discussed above, please work with the messages file to build a comprehensive set of regular expressions to audit the messages in the syslog on this system.

The usual procedure when defining these expressions is to add a few lines, and then run swatch with that config against the file to see what additional expressions should be added. The last line in your configuration is typically a “watchfor” line so that anything not intentionally excluded will be printed.

(Hint: We recommend piping the results of the swatch through the “uniq” command and that output through more:

swatch --config-file=/tmp/swatchconfig --examine=/var/log/messages | uniq | more

Exercise 3: Password Assessment

Purpose: Demonstrate the simplicity of cracking passwords under Unix and demonstrate the value of a strong password policy or additional authentication mechanisms. We will use a simple online cracker to accomplish this task. It is important to note that there are plenty of offline crackers, so this could be performed offline as well, but we might not want to bear the responsibility of carrying the password file off site.

Using the “John” tool, we will assess the strength of the passwords on the Linux system. We already have the root password for this system, so it is the natural place to start. On many modern Unix systems, the password hashes are stored in `/etc/shadow`, which is readable only by root. This is because the previous place they were stored, `/etc/passwd`, must be world readable in order for much of the system to function. Because Unix password hashes are so trivially cracked these days, hiding them from prying eyes became important.

This exercise also demonstrates how we can work with a system administrator to conduct the password audit right on the live system. In this case, we’ve already done the legwork of downloading and installing the software. In the real world, this should take the system administrator less than ten minutes to complete. We could even bring the program with us on some removable media. One thing we definitely want to bring along is the dictionary file that we will use for cracking. In this case, we will use the standard Unix dictionary and the dictionary that is supplied with John. There are some fantastic common password dictionaries out on the Internet as well. I’d recommend getting one of these and using it regularly.

To perform the assessment, boot your VMware “Unix Web Server” system if you have not done so already. Earlier today, you should have created a directory in the /tmp tree with your first initial and last name. Please change into this directory before proceeding.

The tool that we want to use to do the password assessment is “john.” Try running the tool with no options and let’s see what we find out.

As you can see, there are a lot of possible options that we can use when running “john”. In fact, this is one of the best Unix-based password assessment tools available. Even though there are so many options, the tool is actually extremely easy to use!

To crack passwords on this system, which password file do we need to point “john” at? (Hint: Passwords are “normally” stored in /etc/passwd.)

*(Hint: Under what conditions are the passwords **not** stored in /etc/passwd?)*

(Hint: NIS, NIS+ and Shadow password systems are perfect examples of when the passwords are not stored in /etc/passwd.)

(Hint: Is there evidence on this system that any of these are in use?)

(Hint: "ls /etc/shadow")

Let's run "John" against the password file. Please try this command: "john /etc/shadow".

How quickly does "John" begin to return results?

An interesting feature is that if you interrupt "John", whenever you rerun the tool with the same password file, it will pick right back up where it left off. The current results are stored in your current directory as a "john.pot" file (already recovered passwords) and a restore file, which is used to resume.

John should still be running. If it isn't, please restart it. While it is running, press any key. What happens?

So John also allows us to see current statistics so that we can see how far along it is in the password-cracking process. We would expect John to take a good amount of time to break the password on our system, so you might want to interrupt the system rather than waiting.

Section 4

Exercise 1 : Building a Tools CD

Purpose: To demonstrate the process used to create a CD of statically compiled, dynamically compiled and supporting libraries for use during the audit of a Unix system.

The process demonstrated in this exercise can be used on any Unix system to create a usable CD of both audit and forensic tools. The very last exercise for today will demonstrate the use of a similarly created CD.

To get started, we first need to boot our VMware “Unix Web Server” system and log in as root.



For this exercise, we will create our CD in the /tmp directory tree. Please use the “cd” utility to switch to “/tmp” and create a directory there bearing your first initial and last name. Once this has been created, please change into that directory. (If you have already created this directory, there is no need to repeat this step unless you have rebooted the Knoppix system.)

Before we go any further, we need to assemble a list of tools that we want on the CD. Please write the names of the various tools that you would like on a CD below:

Here is a sample list of tools:

/bin/sh	/bin/tar	/bin/mkdir
/bin/cat	/bin/chown	/bin/ps
/bin/ls	/bin/chmod	/bin/mv
/usr/bin/script	/bin/chgrp	/bin/pwd
/usr/sbin/lsof	/bin/cp	/bin/rm
/usr/sbin/tripwire	/bin/more	/bin/rmdir
/usr/sbin/twadmin	/bin/mount	/bin/su
/usr/sbin/twprint	/bin/df	/bin/touch
/bin/netstat	/usr/bin/du	/bin/uname
/sbin/ifconfig	/bin/fuser	/bin/umount
/sbin/fdisk	/bin/echo	/usr/bin/diff
/sbin/lsmode	/bin/grep	
/bin/dd	/bin/gunzip	

Now that we have a list of tools, the next step is to begin copying them into the CD directories. To do this, you should already be in “/tmp/dhoelzer” or whatever would correspond to your first initial and last name. Within this directory, let’s make two more directories.

Please use the “mkdir” command to create the “bin” and “lib” directories.

The “bin” directory is where we are going to put all of the programs that we want on the CD. The “lib” directory will be used for any shared libraries that prove to be necessary. Let’s get started.

One by one, copy each of the files in your list of files for your CD to the “/tmp/yourname/bin” directory. It should look something like this:

```

C:\WINNT\System32\cmd.exe - telnet linux.sans.org
linux:/tmp# mkdir dhoelzer
linux:/tmp# cd dhoelzer
linux:/tmp/dhoelzer# mkdir bin
linux:/tmp/dhoelzer# cp /bin/sh bin
linux:/tmp/dhoelzer# cp /bin/cat bin
linux:/tmp/dhoelzer# cp /bin/echo bin
linux:/tmp/dhoelzer# cp /bin/ls bin
linux:/tmp/dhoelzer#

```

After you copy all of these files into the appropriate directory, the next step is to identify required shared libraries and copy those to the "lib" directory. To do so, follow the example in this screenshot:

```

C:\WINNT\System32\cmd.exe - telnet linux.sans.org
linux:/tmp/dhoelzer# ldd bin/sh
        libc.so.5 => /lib/libc.so.5 <0x40017000>
        libdl.so.2 => /lib/libdl.so.2 <0x40055000>
        libc.so.6 => /lib/libc.so.6 <0x40059000>
        /lib/ld-linux.so.2 => /lib/ld-linux.so.2 <0x40000000>
linux:/tmp/dhoelzer# cp /lib/libc.so.5 lib
linux:/tmp/dhoelzer# cp /lib/libdl.so.2 lib
linux:/tmp/dhoelzer# cp /lib/libc.so.6 lib
linux:/tmp/dhoelzer#
linux:/tmp/dhoelzer# ldd bin/cat
        libc.so.6 => /lib/libc.so.6 <0x40017000>
        /lib/ld-linux.so.2 => /lib/ld-linux.so.2 <0x40000000>
linux:/tmp/dhoelzer# cp /lib/ld-linux.so.2 lib
linux:/tmp/dhoelzer# _

```

Please note that because we copied the /lib/libc.so.6 file when copying the libraries over for /bin/sh, there is no need to copy a new version when we do /bin/cat.

We included a script on the machine that is pretty handy for this. If you look at /bin/copy_libs, this is a Perl script that will automatically copy the proper libraries for you! To use it, from /tmp/yourname, run this command :

```
ls bin/* | copy_libs
```

This will automatically run "ldd" against every binary program and copy the appropriate library into the lib directory. Below is a screenshot of this tool:

```

C:\WINNT\System32\cmd.exe - telnet linux.sans.org
linux:/tmp/dhoelzer# cat /bin/copy_libs
#!/usr/bin/perl

foreach $i < <STDIN> >
{
    chomp($i);
    open(COMMAND, "/usr/bin/ldd $i !");
    foreach <<COMMAND>>
    {
        (<$marker, $pointer, $lib, $junk> = split(/ /);
        if($lib ne "dynamic") {
            print "Copying $lib for binary $i\n";
            system("cp $lib lib");
        }
    }
    close(COMMAND);
}

linux:/tmp/dhoelzer# _

```

Now that we copied all of the files into the proper locations, we are ready to make the disk. To do this, we need to figure out how big the disk image needs to be. To find this

out, we'll use the "du" utility to figure out the disk usage. From the /tmp/yourname directory, run "du -s *".

How much disk space is needed for your CD?

```
C:\WINNT\System32\cmd.exe - telnet linux.sans.org
linux:/tmp/dhoe1zer# du -s *
12016   bin
1720    lib
linux:/tmp/dhoe1zer#
```

Looking at the results that we come up with in the example, it looks like we need around 14 megabytes to store this data. This is actually pretty small! Really, we could store the utilities on a business card-sized CD-R. Now that we know

what the resource requirements are, we can create the actual image. To do this, perform the following steps:

1. Use the "dd" utility to create a blank image file. The command to run is:
dd if=/dev/zero of=CD.img bs=1024 count=15000
2. Now that we created the file to store the image in, we need to format it. To do this, we will use the "mkfs" utility. Enter the following at the command prompt:
mkfs CD.img
The system will ask you whether you are sure. Answer "y".
3. Use the "mkdir" command to create a temporary mount point:
mkdir mnt
4. Mount the new filesystem on the temporary mount point:
mount -o loop CD.img mnt
5. Copy the Binary and Library directories onto the CD:
cp -r bin mnt
cp -r lib mnt
6. Unmount the CD image:
umount mnt
7. Test mounting the volume read only to simulate burning the image onto a CD:
mount -o ro -o loop CD.img mnt

Voila! If you "cd" into the mnt directory, you will see your bin and lib directories, and within those will be all of the utilities and libraries! If you move back to /tmp/yourname, you can "umount mnt" to unmount the image. At this point, you have an image that you can burn to a CD or any other media and mount up later for use!

Exercise 2: A Unix Conformance Audit

Purpose: To give you the opportunity to compare an existing baseline from a Unix system to the running system and look for variances. A portion of the exercise will require you to consider whether or not the variances observed are exceptions or expected.

Introduction

An auditor has previously performed a variety of baseline functions against the VMware system that we are using. We have been asked to perform a manual verification to see whether the machine has changed. The previous auditor has provided us with a simple step-by-step guide on how the baseline was generated so that we can find differences easily.

Action List 1

Purpose: Log in and prepare a temporary space for the audit.

1. **Log in to the host being audited.**
(For the audit, we are using a VMware system, but you would perform essentially these exact steps on a real Unix server.)
2. **Become root.**
(Once logged in as the auditor, type *“sudo su”*)
3. **“cd” to /tmp**
4. **“mkdir auditorname-results”**
(At this point, you should create a directory using the following convention and your own name: *“dhoelzer-results,”* using the first initial of your first name followed by your last name.)
5. **“cd auditorname-results”**

Action List 2

Purpose: Gather data about the system that will be compared to the existing baseline. (In this section, first try to figure out how to create the piece of information required. The correct answers follow in the next section, which is another copy of “Action List 2.”)

1. **Create a text file listing all listening TCP ports and the corresponding processes.**
2. **Create a text file listing all listening UDP ports and the corresponding processes.**
3. **Using another tool, create another text file listing all listening network ports without process information.**
4. **Create a text file listing all running processes where the parent process is process 0 or process 1.**
5. **Create a text file listing all files on the system that are either set UID or set GID.**

Action List 2

Purpose: Gather data about the system that will be compared to the existing baseline.

1. **Create a text file listing all listening TCP ports and the corresponding processes.**
“lsof -i | grep LISTEN > lsof.tcp.observed”
2. **Create a text file listing all listening UDP ports and the corresponding processes.**
“lsof -i | grep UDP > lsof.udp.observed”
3. Using another tool, create another text file listing all listening network ports without process information.
“netstat -an | grep LISTEN > netstat.observed”
4. Create a text file listing all running processes where the parent process is process 0, process 1, or process ID 2. Upstart systems, you will notice, start services slightly differently, meaning that startup processes are owned by kthreadd rather than init.
ps -eaf | grep “[a-zA-Z]\+ \+[0-9]\+ \+[012] \+ > processes.observed”
5. **Create a text file listing all files on the system that are either set UID or set GID.**
“find / -perm +06000 -fls suid.sgid.observed”

Aside:

Please describe what each of these commands do:

```
lsof -i | grep LISTEN > lsof.observed
```

```
lsof -i | grep UDP > lsof.udp.observed
```

```
netstat -an | grep LISTEN > netstat.observed
```

```
ps -eaf | grep '[a-zA-Z]\+ \+[0-9]\+ \+[012] \+ >  
processes.observed
```

```
find / -perm +06000 -fls suid.sgid.observed
```

Answers:

Please describe what each of these commands do:

```
lsof -i | grep LISTEN > lsof.tcp.observed
```

This command will produce a listing of all network ports currently held open by any process, and then use “grep” to reduce this list down to ports that are in state LISTEN. This will find all listening TCP services and which process owns them. The output is redirected to a file called “lsof.observed.”

```
lsof -i | grep UDP > lsof.udp.observed
```

This command will produce a listing of all network ports currently in an open state. This output is sent through “grep” and reduced to only ports bound to protocol UDP, thus producing a list of UDP services and the processes that own them. The output is redirected to a file called “lsof.udp.observed.”

```
netstat -an | grep LISTEN > netstat.observed
```

This command will produce a list of all active network ports without resolving network or service names. This list is sent through “grep” to reduce the list to only ports in state LISTEN, thus ignoring active or transient connections. The output is redirected to a file called “netstat.observed.”

```
ps -eaf | grep “[a-zA-Z]\+ \+[0-9]\+ \+[012] \+” > processes.observed
```

This command will produce a listing of all processes currently running on the system with the parent process ID included. The output is filtered through “grep” using a regular expression to identify processes with a parent process id of zero or one ([01]). These processes are generally started at boot time by the init process or as pseudo processes by the kernel itself. The output is redirected to a file named “processes.observed.”

```
find / -perm +06000 -fls suid.sgid.observed
```

This command will traverse the file system from the root directory identifying any files that have either the set UID or the set GID bit set. The “-fls” option redirects a copy of the “ls -l” equivalent output to the specified file, which is “suid.sgid.observer.”

Action List 3

Purpose: Compare the observed results with the stored baselines. For any observed differences, we should be able to distinguish between actual exceptions and normal variance.

1. **Using the following command, make a copy of the stored baselines for comparison.**
`"cp /var/adm/baselines/* ./"`
(This will copy all of the files in "/var/adm/baselines" into the current directory, which should be the temporary audit directory that we created in action list 1.)
2. **For each of the observed files, identify the corresponding baseline file and compare them.**
(Hint: "diff lsof.tcp lsof.tcp.observed" will show us all of the differences between the two files.)
3. **For each observed difference, please identify whether or not it is an exception or expected/acceptable variance. Please list all exceptions and variances below in the appropriate sections.**

Exceptions:

Please list and detail all exceptions found with possible causes.

Example:

```
linux:/tmp/dhoelzer # diff lsof lsof.observed
1,9c1,8
< inetd 545 root 4u IPv4 925 TCP *:time (LISTEN)
< inetd 545 root 6u IPv4 927 TCP *:telnet (LISTEN)
< inetd 545 root 7u IPv4 928 TCP *:login (LISTEN)
< inetd 545 root 8u IPv4 929 TCP *:finger (LISTEN)
< portmap 557 bin 4u IPv4 1029 TCP *:sunrpc (LISTEN)
< sshd 611 root 3u IPv6 1328 TCP *:ssh (LISTEN)
< sendmail 685 root 3u IPv4 1691 TCP localhost:smtp (LISTEN)
< httpd 853 root 16u IPv4 2950 TCP *:http (LISTEN)
< httpd 854 wwwrun 16u IPv4 2950 TCP *:http (LISTEN)
---
> inetd 551 root 4u IPv4 1054 TCP *:time (LISTEN)
> inetd 551 root 6u IPv4 1056 TCP *:telnet (LISTEN)
> inetd 551 root 7u IPv4 1057 TCP *:login (LISTEN)
> inetd 551 root 8u IPv4 1058 TCP *:finger (LISTEN)
> portmap 563 root 4u IPv4 1252 TCP *:sunrpc (LISTEN)
> sendmail 674 root 3u IPv4 1805 TCP localhost:smtp (LISTEN)
> httpd 835 root 16u IPv4 2864 TCP *:http (LISTEN)
> httpd 869 root 16u IPv4 2864 TCP *:http (LISTEN)
```

The "diff" command reveals that the "sshd" daemon was running and has been disabled. There is no change control documentation available to demonstrate why this service would have been disabled.

Expected or Acceptable Variances:

Please list all expected or acceptable variances below and explain why each is expected.

Example:

```
linux:/tmp/dhoelzer # diff lsof lsof.observed
1,9c1,8
< inetd  545  root  4u  IPv4  925   TCP *:time (LISTEN)
< inetd  545  root  6u  IPv4  927   TCP *:telnet (LISTEN)
< inetd  545  root  7u  IPv4  928   TCP *:login (LISTEN)
< inetd  545  root  8u  IPv4  929   TCP *:finger (LISTEN)
< portmap 557  bin   4u  IPv4  1029  TCP *:sunrpc (LISTEN)
< sshd   611  root  3u  IPv6  1328  TCP *:ssh (LISTEN)
< sendmail 685  root  3u  IPv4  1691  TCP localhost:smtp (LISTEN)
< httpd  853  root  16u IPv4  2950  TCP *:http (LISTEN)
< httpd  854  wwwrun 16u IPv4  2950  TCP *:http (LISTEN)
---
> inetd  551  root  4u  IPv4  1054  TCP *:time (LISTEN)
> inetd  551  root  6u  IPv4  1056  TCP *:telnet (LISTEN)
> inetd  551  root  7u  IPv4  1057  TCP *:login (LISTEN)
> inetd  551  root  8u  IPv4  1058  TCP *:finger (LISTEN)
> portmap 563  root  4u  IPv4  1252  TCP *:sunrpc (LISTEN)
> sendmail 674  root  3u  IPv4  1805  TCP localhost:smtp (LISTEN)
> httpd  835  root  16u IPv4  2864  TCP *:http (LISTEN)
> httpd  869  root  16u IPv4  2864  TCP *:http (LISTEN)
```

Although “diff” reports that there are many differences, with the exception of the removal of the “sshd” daemon, which is handled as an exception, the differences are the process ID numbers. It is not unreasonable to find that processes have slightly different process ids from boot to boot; therefore, these are expected variances.

Action List 4:

Purpose: Examine the system with a file integrity checking tool to find any modified files. We will store the output into a text file for easy reference.

1. **Run “tripwire” against the system in “checking” mode to find any modified files.**

`“tripwire -m c > tripwire.report”`

(This will check the filesystem against the existing database and redirect the output into a file called “tripwire.report.”)

2. **Evaluate the “Tripwire” output for variances and exceptions.**

After reading the report, please express your thoughts regarding the condition of this system:

Day 6

NetWars: Audit the Flag!

Welcome to Audit the Flag using the NetWars simulation engine! Following a brief introduction and explanation from your instructor, you will be provided with an access code that will allow you to join the game. We invite you to sign up and take your time in this guided self-paced exploration of a simulated enterprise network.

You will find that the vast majority of the questions posed can be answered using both the systems provided in the simulation (of course) and the relevant course book from this week. As you progress through the simulation, you are being asked to answer questions pertaining to the configuration of the systems being examined. However, we would ask that you keep an eye on the metaconfiguration level.

Possibly periodically but definitely at the end of the simulation experience, your instructor will guide a discussion attempting to elicit from you two things. First, possible root causes of the issues found and additional interviews, questions, or inspections that you might want to make. Secondly, the instructor will be looking for process-level solutions to some of the issues discovered. Feel free to participate and to be creative! There really are no wrong answers!

If you run into trouble during the simulation, here are a few suggestions:

1. Don't be afraid to use the hints. The hints do not cost you points (while guessing does).
2. Don't be shy about collaborating with others. Especially if you have expertise in one area and recognize that someone else has strengths in another, leverage that expertise by asking him or her questions. ***Do not ask for his or her answers!*** Although this isn't (currently) a graded exercise, learning occurs only during the search for knowledge and understanding, not when keying in someone else's answer!
3. Don't be afraid to call the instructor or a TA over for assistance! Also, don't be frustrated with any answers that he or she might give you. Rather than giving you "the answer," these educators will instead try to guide you toward an answer through questions. An exception to this policy, of course, is "Hey, this doesn't work... I can't connect!" In these cases, we'll do whatever we can to help you to resolve the issue.

Appendices

Selected Answers

ESXi Issues

- ***What is the potential impact of leaving the datastore unprotected?***

Leaving the datastore easily accessed from a web browser offers an attacker the opportunity to simply download the entire disk image and potentially memory snapshots of the virtual system. If the attacker is able to obtain a copy of the drive, then there is really no need to bother guessing passwords because the data itself has been exposed!
- ***What are some effective methods for protecting against web-based datastore browsing?***

The two simplest adjustments that can be made are to, first, select a long passphrase for any user whose account has the rights necessary to browse the datastore. Secondly, all VMKernel interfaces should be well protected. The web-based interface is accessible only on VMKernel interfaces, none of which should be easily accessible from the intranet, much less the Internet!
- ***The following questions can be answered using the Summary tab:***
 - At least one configuration issue is being brought to your attention. What are these issues?
 - There is no password for the root user.
 - What is the current licensing status of this system?
 - Evaluation Mode
 - How many processor sockets are installed?
 - 2
 - How many NICs are installed?
 - 1
 - Are there any potential findings concerning the number of NICs installed?
 - Yes. Having only one physical adapter creates several issues. First, it is quite likely that VMKernel management will be occurring over this interface, which is being shared for other accessible services. VMKernel ports should be isolated. In addition, having only one physical NIC will likely cause availability problems with more than a few virtual machines running on this host. It would be best to have a ratio as close to 1:1 as possible when it comes to VMs to NICs, or to at least have performed utilization calculations demonstrating that the bandwidth available is sufficient for the requirements of the systems.
 - How many networks are configured on this ESXi system?
 - Three: VM Network, Services, and Intranet
- ***The following questions can be answered using the Configuration tab:***
 - How much memory is installed on this virtual server?
 - 2047.4 MB

SANS Advanced Systems Audit Workbook

- What is the total capacity of the storage available?
 - 3.00 GB
 - What is the name of the installed datastore?
 - ESXi Internal Datastore
 - Is the installed datastore an SSD drive?
 - No
 - Which file system is in use on the datastore?
 - VMFS5
 - Which networks are connected to the physical network adapter?
 - VM Network and the Management Network
 - What is the VLAN ID of the “Services” network?
 - 100
 - Is IPV6 enabled on vmnic0?
 - Yes.
 - To which switch is vmnic0 attached?
 - vSwitch0
 - Is there anything wrong with the time configuration?
 - Yes. The NTP Client is stopped and there are no servers configured. Without this configured, the time will drift leading to inaccuracies in the log and potential authentication issues. This can potentially affect all installed virtual systems as well.
 - Is the system configured to accept domain authentication?
 - No.
- ***To answer the following questions you must click the “Advanced Settings” option under “Software”:***
 - Are the system logs configured to be sent to a remote system?
 - Yes.
 - What is the IP address of the system?
 - 192.168.2.254
 - What is the current setting for maximum log size that will cause the logs to be rotated?
 - 1024 Kb or 1 Megabyte

Firewall Validation – Optional

This exercise is included for you in the event that you are in a position where you need to actually perform the validation of a firewall yourself. You might also find this information very useful if you need to oversee the validation of a firewall (or other network control) or if you want to share the details of how to do so with a network or security administrator/engineer.

Automated Scanning and Analysis

Certainly, you can see why we need to validate a firewall or router at this point. However, you are also likely thinking, “Wow, that’s going to be an enormous amount of work!” Never fear! If there’s a task that you need to do more than one time, automation is the way to go. We will actually spend time in class later this week learning how to write basic scripts. For now, we will leverage an existing set of scripts.

On the Internal host, please close Wireshark if you have it open, and open a root shell. Once the shell is open, please use the “cd” command to change into the AuditcastsScripts/PCI directory.

```
root@internal:/home/auditor# cd AuditcastsScripts/PCI/
root@internal:/home/auditor/AuditcastsScripts/PCI# ls
cisp_analyze  cisp_scanner  cisp_vuln_scan
cisp_ciphers  cisp_sniffer  sample_pci_report.html
root@internal:/home/auditor/AuditcastsScripts/PCI# █
```

As you can see, there are a number of scripts in this directory. The same scripts are in the same place on the External system. For our class, we will use only three of these: `cisp_scanner`, `cisp_sniffer`, and `cisp_analyze`.

To get started, ***please execute the following command on the Internal host:***

```
./cisp_sniffer
```

When you execute this command, you will be prompted to indicate whether this sniffer is inside or outside the firewall. Obviously, because we are in the Internal host, you should select “i” for internal. You will then be prompted to select the Ethernet adapter to listen on. Just press the Enter key to accept the default.

SANS Advanced Systems Audit Workbook

```
root@internal:/home/auditor/AuditcastsScripts/PCI# ./cisp_sniffer
Copyright (C) 2005, All Rights Reserved -- Cyber-Defense.Org
David Hoelzer
http://www.cyber-defense.org
-----
```

PCI CISP Sniffer driver

Is this the internal sniffer or the external sniffer? [I/E] i
Which interface should I listen on? [eth0]

Starting sniffer on eth0. The log will be stored in
/home/auditor/AuditcastsScripts/PCI as inner_capture
When the scan has been completed, hit control-C to terminate.

```
-----
tcpdump: listening on eth0, link-type EN10MB (Ethernet), capture size 65535 bytes
_
```

Now that the sniffer is running on the inside, let's switch to the external system and start up the scanner. **Please leave the sniffer running and switch to the External host. Once there, use a root command shell to switch into the AuditcastsScripts/PCI directory so that you can access the cisp_scanner script. Please execute the script with the following command:**

```
./cisp_scanner
```

When you run the scanner, you will first be prompted for "Quick mode." Quick mode is only for testing things out. However, because it speeds things up, we'll use this in the lab to simulate the full test. This will scan only the first 80 ports through the firewall. Answer "y" to this question.

```
root@external:/home/auditor/AuditcastsScripts/PCI# ./cisp_scanner
Copyright (C) 2005, All Rights Reserved -- Cyber-Defense.Org
David Hoelzer
http://www.cyber-defense.org
-----
```

The scanner will normally probe all possible ports on the destination network. If you would like to run in 'quick' mode for testing purposes, the scanner will only probe the first 80 ports, thus running approximately 60 times faster.

Run in quick mode? [y/N]y

Next, you will be prompted to enter the address of an existing host behind the firewall. For our purposes, we will use the address of the Internal scanner. This host is located at 10.17.1.20. Enter this value and press Return.

SANS Advanced Systems Audit Workbook

The scanner needs information about the target network or systems. If you are scanning from inside to outside, any external address will do as a target, but we recommend you choose something that is unallocated or blocked at your edge router to prevent the packets from hitting another party. If you are scanning inbound, the target should be a real host IP address behind the firewall.

Enter a target address: 10.17.1.20

The next prompt will ask you for the address of the firewall. The answer to this question should be the address of the interface that is closest to the scanner. You should have already identified this address earlier in this lab. Fill in that information here and press the Return key.

The next prompt asks whether you are working in Internal or External mode. The answer to this question is determined by the position of the scanner. Because the scanner is currently external, the answer will be "e". Select this and press Return.

The scanner needs to know if it should operate in internal or external mode since the scans performed are slightly different.

Operate in internal or external mode? [I/e]e

The scanner will now automatically determine the MAC address of the firewall. This is done to accomplish what we did when we manually configured a route. This allows the scanner to deliver packets for internal hosts directly to the firewall interface to see whether they are passed in.

The last question configures a real IP address to use from the outside. The scanner will spoof packets from this source (in addition to localhost and others) in an effort to bypass the firewall. Let's leave this address as-is and press Return.

In external mode, the scanner will need to know the ethernet address of the firewall. The scanner will automatically ping the firewall in an effort to ascertain the MAC address. It is critical that this scanner be located on the same physical segment as the firewall's external interface for this to function properly.

The MAC of the firewall is 00:0c:29:d1:15:ae

The scanner must use some real public IP address for the scans from the outside. The scanner uses 5.1.1.1 by default since this is an unallocated address. Is this acceptable?

Use 5.1.1.1 as the public address? [Y/n]

At this point, the system will begin scanning the firewall. Simply wait until all of the inbound scanning tests are completed!

SANS Advanced Systems Audit Workbook

```
Performing external scan of firewall address.
Initial external scan completed.
Beginning fragmentation and flag scanning through firewall.
Finished complete loop for TCP option "-fS".
Finished complete loop for TCP option "-fS -fA".
Finished complete loop for TCP option "-fS -fF".
Finished complete loop for TCP option "-f-".
TCP testing completed
Beginning ICMP sweep
ICMP testing complete
UDP testing begins
Performing an NMap scan of the firewall, 192.168.129.128...
```

Scan complete!

```
Probing Completed!root@external:/home/auditor/AuditcastsScripts/PCI# █
```

When you receive the message that the probing has been completed, it's time to switch things around and scan in the other direction. ***Please execute the "cisp_sniffer" tool on the external system. Configure it as the external sniffer.***

After you have the sniffer running on the External system, switch to your Internal host and stop the sniffer by pressing CTRL-C.

```
Starting sniffer on eth0. The log will be stored in
/home/auditor/AuditcastsScripts/PCI as inner_capture
When the scan has been completed, hit control-C to terminate.
```

```
-----
tcpdump: listening on eth0, link-type EN10MB (Ethernet), capture size 65535 bytes
^C192 packets captured
192 packets received by filter
0 packets dropped by kernel
```

```
root@internal:/home/auditor/AuditcastsScripts/PCI# █
```

Now start the ./cisp_scanner tool on the Internal host. Please answer the prompts as follows:

```
Quick mode:   Yes
Target address: Enter the address of your External system
Firewall address: 10.17.1.1
Mode:         i
```

The scanner will now start. Please wait patiently for it to finish.

When the scanner finishes running, once again switch to the External system and stop the sniffer. With the sniffer stopped, please enter the following command on the External system:

```
scp auditor@10.17.1.1:/home/auditor/AuditcastsScripts/PCI/i* ./
```

Note that you are entering this command from within the root prompt. In addition, you are still located in the AuditcastsScripts/PCI directory.

When you execute this copy command, you will be prompted to accept the “host key” for the remote server. Please just answer “yes.” You will also be prompted for the password for the Auditor account. The password is “Password1.”

```
root@external:/home/auditor/AuditcastsScripts/PCI# scp auditor@10.17.1.1:/home/a
uditor/AuditcastsScripts/PCI/i* ./
The authenticity of host '10.17.1.1 (10.17.1.1)' can't be established.
ECDSA key fingerprint is 24:0e:58:31:32:88:ff:b2:ac:80:6b:0e:c8:a2:03:c9.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '10.17.1.1' (ECDSA) to the list of known hosts.
auditor@10.17.1.1's password:
Permission denied, please try again.
auditor@10.17.1.1's password:
i_nmap_scan          100%  409    0.4KB/s   00:00
inner_capture        100%  14KB  14.3KB/s   00:00
i_scan_settings      100%   37    0.0KB/s   00:00
root@external:/home/auditor/AuditcastsScripts/PCI# █
```

This command will copy all of the files that were generated during the Internal portion of the scanning and relocate them to the External host. With that done, we can move on to the analysis phase!

To perform the analysis, please type the following command on the External host:

```
./cisp_analyze
```

This command will automatically discover which tests have actually been run and then begin prompting you for information about what it has found. *The information for answering these questions can be found on page 49.*

There are two kinds of questions that you will be asked. First, are there ports that the **Firewall** requires to be open on the inside or outside? In this case, because it is a NAT, we would expect any inbound services that are needed to be listening on the firewall itself.

The second type of question is about the **Infrastructure**. These define ports that must pass *through* the firewall, though the firewall itself might not be listening on them.

SANS Advanced Systems Audit Workbook

```
Please enter the port number: 80
Are there more required ports? [y/N]
Added 1 permitted inbound ports on the inside of the firewall.
Are there ports that should be open on the outside of the firewall? [y/N]y
Please enter the port number: 25
Are there more required ports? [y/N]y
Please enter the port number: 53
Are there more required ports? [y/N]
Added 2 permitted inbound ports on the outside of the firewall.
Does your infrastructure require more outbound ports? [y/N]y
To enter a required port, you must first specify the
protocol type. Is this port a TCP or a UDP port? [T/u]t
tcp selected. Please enter the outbound port number: 25
tcp/25 added.
Does your infrastructure require more outbound ports? [y/N]y
To enter a required port, you must first specify the
protocol type. Is this port a TCP or a UDP port? [T/u]u
udp selected. Please enter the outbound port number: 53
udp/53 added.
Does your infrastructure require more outbound ports? [y/N]
reading from file inner_capture, link-type EN10MB (Ethernet)
reading from file inner_capture, link-type EN10MB (Ethernet)
reading from file outer_capture, link-type EN10MB (Ethernet)
reading from file outer_capture, link-type EN10MB (Ethernet)
Report completed. View 'pci_report.html' for the results.

root@external:/home/auditor/AuditcastsScripts/PCI# █
```

After answering these questions, we're ready to see the results. The report has been stored in a file called "pci_report.html." ***To view the report, please type the following on the External system:***

```
firefox pci_report.html
```

You should see the Firefox web browser open with the report displayed!

Please examine this report and see whether there were any unexpected ports open on the firewall or through the firewall.



WMIC Cheatsheet

These are current WMIC facilities available in Server 2012. Not all aliases will be available on your system or a workstation!

WMIC [global_switches] <alias> [options] [format]

Interactive mode:
WMIC

Aliases

ALIAS	Access to the aliases available on the local system
BASEBOARD	Base board (also known as a motherboard or system board) management
BIOS	Basic input/output services (BIOS) management
BOOTCONFIG	Boot configuration management
CDROM	CD-ROM management
COMPUTERSYSTEM	Computer system management
CPU	CPU management
CSPRODUCT	Computer system product information from SMBIOS
DATAFILE	DataFile management
DCOMAPP	DCOM Application management
DESKTOP	User's Desktop management
DESKTOPMONITOR	Desktop Monitor management
DEVICEMEMORYADDRESS	Device memory addresses management
DISKDRIVE	Physical disk drive management
DISKQUOTA	Disk space usage for NTFS volumes
DMACHANNEL	Direct memory access (DMA) channel management
ENVIRONMENT	System environment settings management
FSDIR	Filesystem directory entry management
GROUP	Group account management
IDECONTROLLER	IDE Controller management
IRQ	Interrupt request line (IRQ) management
JOB	Provides access to the jobs scheduled using the schedule service
LOADORDER	Management of system services that define execution dependencies
LOGICALDISK	Local storage device management
LOGON	LOGON Sessions
MEMCACHE	Cache memory management
MEMORYCHIP	Memory chip information
MEMPHYSICAL	Computer system's physical memory management
NETCLIENT	Network Client management

SANS Advanced Systems Audit Workbook

NETLOGIN	Network login information (of a particular user) management
NETPROTOCOL	Protocols (and their network characteristics) management
NETUSE	Active network connection management
NIC	Network Interface Controller (NIC) management
NICCONFIG	Network adapter management
NTDOMAIN	NT Domain management
NTEVENT	Entries in the NT Event Log
NTEVENTLOG	NT eventlog file management
ONBOARDDEVICE	Management of common adapter devices built into the motherboard (system board)
OS	Installed Operating System/s management
PAGEFILE	Virtual memory file swapping management
PAGEFILESET	Page file settings management
PARTITION	Management of partitioned areas of a physical disk
PORT	I/O port management
PORTCONNECTOR	Physical connection ports management
PRINTER	Printer device management
PRINTERCONFIG	Printer device configuration management
PRINTJOB	Print job management
PROCESS	Process management
PRODUCT	Installation package task management
QFE	Quick Fix Engineering
QUOTASETTING	Setting information for disk quotas on a volume
RDACCOUNT	Remote Desktop connection permission management.
RDNIC	Remote Desktop connection management on a specific network adapter
RDPERMISSIONS	Permissions to a specific Remote Desktop connection
RDTOGGLE	Turning Remote Desktop listener on or off remotely
RECOVEROS	Information that will be gathered from memory when the operating system fails
REGISTRY	Computer system registry management
SCSICONTROLLER	SCSI Controller management
SERVER	Server information management
SERVICE	Service application management
SHADOWCOPY	Shadow copy management
SHADOWSTORAGE	Shadow copy storage area management
SHARE	Shared resource management
SOFTWAREELEMENT	Management of the elements of a software product installed on a system
SOFTWAREFEATURE	Management of software product subsets of SoftwareElement
SOUNDDEV	Sound Device management

SANS Advanced Systems Audit Workbook

STARTUP	Management of commands that run automatically when users log onto the computer system
SYSACCOUNT	System account management
SYSDRIVER	Management of the system driver for a base service
SYSTEMENCLOSURE	Physical system enclosure management
SYSTEMSLOT	Management of physical connection points including ports, slots and peripherals, and proprietary connections points
TAPEDRIVE	Tape drive management
TEMPERATURE	Data management of a temperature sensor (electronic thermometer)
TIMEZONE	Time zone data management
UPS	Uninterruptible power supply (UPS) management
USERACCOUNT	User account management
VOLTAGE	Voltage sensor (electronic voltmeter) data management
VOLUME	Local storage volume management
VOLUMEQUOTASETTING	Associates the disk quota setting with a specific disk volume
VOLUMEUSERQUOTA	Per user storage volume quota management
WMISET	WMI service operational parameters management

Useful DSQuery Formulae

All users in the domain:

```
dsquery * domainroot -filter "(&(objectclass=user)(objectcategory=person))" -attr  
sAMAccountName  
dsquery user
```

All domain controllers in the forest:

```
dsquery * cn=configuration,dc="domain",dc="org/com/net/local/etc" -filter  
"objectclass=ntdsdsa"  
dsquery server -forest
```

Find Global Catalog servers:

```
dsquery server -forest -isgc
```

Find FSMO Roles:

```
dsquery server -hasfsmo rid  
dsquery server -hasfsmo pdc  
dsquery server -hasfsmo name  
dsquery server -hasfsmo schema  
dsquery server -hasfsmo infr
```

Enumerate OUs:

```
dsquery ou domainroot
```

Locked out users/Users who can't lock out - Not necessarily locked out right now - locked out at some point and not yet logged in:

```
dsquery * domainroot -filter  
"(&(objectcategory=person)(objectclass=user)(lockoutTime=*))"
```

Disabled Users:

```
dsquery user -disabled
```

Get group memberships:

```
dsquery user -name David* | dsget user -memberof -expand
```

Get groups:

```
dsquery group
```

Find members of groups:

```
dsget group "Group DN" -members -expand  
dsquery group | dsget group -members -expand
```

Find computers in the domain:

```
dsquery computer
```

Unused computers:

dsquery computer -inactive <num weeks>¹

Enumerate all group policies in domain:

dsquery * domainroot -filter (objectcategory=grouppolicycontainer) -attr displayname

¹ Requires that the Active Directory is at a 2003 functional level or higher only

User Account Control Bit Values

Description	Hex Value	Decimal Value
Logon script will execute	0x01	1
Account is disabled	0x02	2
User home directory must exist	0x08	8
This account is locked out	0x10	16
User is not required to have a password	0x20	32
User cannot change password	0x40	64
An encrypted password (rather than typical hashing) can be sent	0x80	128
This is a local user account that allows a user in an untrusted domain to access resources in this domain	0x100	256
Normal user account	0x200	512
This account creates an interdomain trust	0x800	2048
This account is a Workstation or Server account	0x1000	4096
This account is the backup domain controller for this domain	0x2000	8192
This user's password does not expire	0x10000	65536
This is a Majority Node Set account	0x20000	131072
This account requires a smartcard to authenticate	0x40000	262144
This account is trusted for Kerberos delegation	0x80000	525288
This user might not be delegated for Kerberos even if a service account is trusted for delegation	0x100000	1048576
The user security principal may only use DES keys for encryption	0x200000	2097152
Kerberos preauthentication is not required for logon	0x400000	4194304

SANS Advanced Systems Audit Workbook

The password for this account has expired	0x800000	8388608
This account is trusted for delegation (might impersonate another account)	0x1000000	16777216
This account represents a read-only domain controller (RODC)	0x40000000	67108864

Domain controller accounts are typically set to 0x82000 (532480) because they are both member servers and trusted for delegation.

This information based on multiple Microsoft references, including [http://msdn.microsoft.com/en-us/library/ms680832\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/ms680832(v=vs.85).aspx)

To perform bitwise testing, the following two rule OIDs are available:

- 1.2.840.113556.1.4.803 Bitwise AND to check for *all* specified bits
- 1.2.840.113556.1.4.804 Bitwise OR to check for *any* of the specified bits

Sample Scripts

Unix Audit Script

The following script is an example of what you might use to initially collect data, periodically check for changes in configuration or even automatically monitor for change on a Unix system. You are free to use this script under the terms of the license explained within the source code itself.

```
#!/bin/sh

COMMENT=

if [ $COMMENT ] ; then

    Copyright 2007, Cyber-Defense and EnclaveForensics
    All Rights Reserved
    Written by David Hoelzer

    LICENSE
    -----
    You may not redistribute this script in any form. Please direct
    interested
    parties to www.cyber-defense.org if they would like a copy. You
    are free
    to use this script upon the condition that neither this script
    nor any
    part thereof is incorporated into any commercial product or
    tool.

    This is a script to allow you to quickly and repeatedly
    collect baseline audit data from a UNIX based system. If you
    add or
    think of any interesting tests to add to this script please
    forward them
    to us at info@cyber-defense.org so that others can benefit!

    CUSTOMIZATION
    -----
    If you are looking to customize the data collection in this
    script you
    will want to scroll down to the bottom of the file, below the
    functions.

fi

function YesNo ()
{
    RESPONSE=x
    while [ $RESPONSE != "y" ] && [ $RESPONSE != "n" ] ; do
        echo -n " (Yes or No [y/n]) "
        read -n 1 RESPONSE
    done
}

```

SANS Advanced Systems Audit Workbook

```
done
test $RESPONSE == "n"
}

function Separator ()
{
    echo +-----
}

function RunTest ()
{
    TEST=$2
    TESTNAME=$1
    if [ -z "$TEST" ] ; then
        echo Empty test called in RunTest.  Exiting.
        exit 50;
    fi
    if [ -z "$TESTNAME" ] ; then
        echo Empty test name called in RunTest for $TEST.
        Exiting.
        exit 51;
    fi
    Separator >> $OutputFile
    echo "| $TESTNAME" >> $OutputFile
    Separator >> $OutputFile

    # There are difficulties getting certain commands with pipes to execute
    # within the shell.  A simple kludge is to shove the commands into a
    # file
    # and execute the file.
    echo $TEST > __trunme
    { /bin/sh __trunme; } >> $OutputFile
    rm -f __trunme
    if [ $? -ne 0 ] ; then
        echo "Error running $TESTNAME: $TEST"
        echo "Bailing out."
        exit 5
    fi
}

function Header ()
{
    Separator > $OutputFile
    echo " `basename $0` test for `hostname` by `whoami` on `date`"
    >> $OutputFile
    Separator >> $OutputFile
}

function GetOutputFile ()
{
    OutputFile=`hostname`-`date +%s`-`basename $0`
}

function GetRunlevel ()
{
    RunLevel=`awk -F: '/^id/ {print $2;}' /etc/inittab`
}
```

SANS Advanced Systems Audit Workbook

```
}
if [ ! -z $1 ] ; then
    if [ ! -s $1 ] ; then
        echo "You requested a comparison but did not provide a
valid filename."
        exit 3
    fi
    BASELINE=$1
fi

#-----
#
# If you are looking to customize this script, you will likely want to
# begin here.  Each of the tests is handled by the RunTest function.
# RunTest takes two arguments, a brief description and the test itself.
#
#-----

echo Running automatic baseline script for `hostname` as `whoami`.
echo If you are attempting to validate a system, please rerun this
script
echo with the name of a baseline file as a command line argument.

GetOutputFile
GetRunlevel
Header

RunTest "Kernel Type/Machine Information" "uname -a"
RunTest "Physical Memory" "free | awk '/Mem/ {print \$2;}'"
RunTest "Mounted Partitions" "mount"
RunTest "Physical Partition Tables" "fdisk -l /dev/?d?"
RunTest "Critical Directory Inventory" "ls -alR /etc /bin /lib /sbin
/usr/lib /usr/bin /usr/sbin"
RunTest "Network interfaces" "ifconfig -a | awk '/^[a-zA-Z]+/ { print
\$1\" - \"\$5; }'"

    # Inventory Listening Ports:
    # A listening TCP port will have the word "LISTEN" on the line.
A listening UDP port will
    # begin with the letters UDP
RunTest "Inventory Listening Ports" "netstat -an | awk '/(^udp)|LISTEN/
{print \$1\" \"\$4;}'"

RunTest "Current Runlevel" "who -r | awk '{ print \$1\" \"\$2;}'"
RunTest "Init Default Runlevel" "awk -F: '/^id/ {print \$2;}'
/etc/inittab"
RunTest "Find Services Started During Startup" "ps eaxl | awk '/^[0-9][
\t]+[0-9]+[ \t]+[0-9]+[ \t]+1[ \t]+/ {print \$13;}'"

    # Inventory SUID and SGID files:
    # -perm with the + option will identify all objects where any of
the listed permissions
    # are set. 04000 is SUID, 02000 is SGID. '-type f' restricts
the list to files only
RunTest "Find SUID and SGID Files" "find / -perm +06000 -type f -ls"
```

SANS Advanced Systems Audit Workbook

```
RunTest "Current Users" "cat /etc/passwd"

# Inventory Root Users:
# Root users in the passwd file will contain a username followed
by colon followed by at
# least one zero followed by a colon
RunTest "Root Users" "awk \"/^[^:]+:[^:]+:0+:/ {print;}\\" /etc/passwd"

# Inventory Blank Passwords:
# A blank password in the shadow file will be a line that has a
username followed
# by a colon followed by a colon
RunTest "Inventory Blank Passwords" "awk -F: '/^[^:]+:::/ {print \$1;}'
/etc/shadow"

# Inventory Active Accounts:
# Inactive accounts in shadow file will contain a line with a
username
# followed by a colon followed by more than one character other
than a colon
# NOTE: Most awks don't understand {2,} so we require [^:][^:]+.
RunTest "Inventory Active Accounts" "awk -F: '/^[^:]+:([^:][^:]+|:)/
{print \$1;}' /etc/shadow"

RunTest "Groups and Membership" "cat /etc/group"

# Rhosts is just a really bad idea these days. We look for and
inventory
# any .rhost files and hosts.equiv
RunTest "Find Unencrypted Remote Trusts" "find / -name .rhost -name
hosts.equiv -ls"

# If the script was invoked with a baseline filename, check for
changes.

if [ ! -z $BASELINE ] ; then
diff $BASELINE $OutputFile > Exceptions
NUMEXCEPTIONS=`wc -l Exceptions | awk '{print \$1;}'`
if [ $NUMEXCEPTIONS -ne 4 ] ; then
echo "Exceptions detected:"
cat Exceptions
else
rm -f Exceptions
fi
fi
```

Troubleshooting

How do I change the keyboard layout settings inside of the virtual machines?

After logging in, use the “loadkeys” tool to adjust the layout. For instance, to adjust to a Danish keyboard layout, you would type “loadkeys dk.” To change to a Spanish keyboard layout, you would type “loadkeys es.”

Web Application Audit Checklist

Basic Configuration

- Is default content or sample content installed on the web server?
 - If so, why is this necessary? Is any of this material executable code?
- Check that directory indexing has been disabled.
- Compare the configuration of the server with manufacturer or community security recommendations. Have appropriate controls and options been configured?
- Have the server headers been sanitized?
- What security baseline has been implemented on the underlying operating system? Has the security of this configuration been audited?
- Does the network architecture support the security and information flow requirements of the web architecture?
- Use at least one automated tool to evaluate the web site.
 - Are there any issues identified by the tools?
 - What sorts of risks do these issues represent?
 - What controls already exist to mitigate the risks?
 - Are the responsible individuals capable of remediating the risks identified?

Authentication

- Is authentication required? If so:
 - How is authentication accomplished?
 - If Basic authentication is used, is it appropriate for the level of sensitivity for the data?
 - If Basic is used, is SSL required?
 - If forms are used, is the POST method used?
 - If forms are used, is SSL required?
 - If certificates are used, how are certificates controlled?
 - If certificates are used, how is the CRL managed?
 - How are account lockouts handled?
 - Are speed bump lockouts in use?

Session Management

- What session management and tracking technique is in use?
- Basic Authentication
 - Is SSL required at all times after the username and password are requested?
 - Is there a sign-off procedure to force the credentials to be dropped from the browser?
 - Is there some form of tracking in place to identify brute-force password guessing attempts?

- How are account lockouts handled?
- URL Rewriting/Hidden form fields
 - Are the session IDs sufficiently random?
 - Are the session IDs sufficiently large?
 - Is the generation of the session ID based on any aspect of the user or password information?
 - Are session IDs sent over appropriately secure paths? If not, how is this mitigated?
 - How are account lockouts handled?
 - Is there a clear sign-off procedure to expire a session ID?
 - Are session IDs perishable?
- Are the session IDs in the session token of sufficient length for the application?
- Are the session IDs secured appropriately based on cloning detection capabilities?
- What type of session hijacking/cloning detection capabilities exist?
- What actions does the application take when a session violation is detected?
- Are the session IDs sufficiently random?
- Do session IDs expire after some period of time?
- Is a valid session required in all appropriate circumstances? How is this enforced or mediated?

Input

- How is input to the application sanitized?
- Is input sanitized in all cases, even if some cases have less restrictive rules?
- Is sensitive information always sent using a POST rather than a GET?
- How robust is the application when dealing with unexpected or illegal input?

Output

- How are error conditions handled?
- Is it possible to cause the application to generate an unhandled error?
- Is encryption used in all cases where sensitive information is returned?
- Are there any anti-caching techniques in use when sensitive information is returned?
- Are all special characters properly stripped or escaped when returned in a web page?