

0 1 1 0 0 1 0 1 0 1 1 0 1 0 0 1
National Computer Forensics Institute
1 1 0 0 0 0 1 0 1 1 1 0 0 1 1 0
0 1 1 0 1 1 0 1 1 0 0 0 0 1 0 1
0 0 1 0 1 0 1 1 0 0 1 1 0 1 1
NITRO
1 1 1 1 0 1 1 1 1 1 1 1 0 1 1 1
1 0 1 0 1 1 1 0 1 0 1 0 1 1 1 0
1 0 0 1 1 1 0 0 1 1 0 1 1 1 0 1
0 0 1 1 0 0 1 1 1 0 1 1 0 0 1 0
0 1 1 0 1 0 0 0 0 1 1 1 1 1 0 0
1 1 0 1 1 0 0 0 1 1 0 0 1 0 0 0
1 0 1 1 1 0 0 1 1 0 1 0 1 0 0 1
Network Intrusion
0 1 1 0 1 0 1 0 1 1 0 1 0 1 1
Responder Program
0 0 1 0 1 1 0 0 0 1 1 0 1 1 1 0
0 0 0 0 1 1 0 0 0 1 1 0 1 1 1 0
0 1 0 1 1 1 0 1 1 0 0 1 1 0 0 1
Volume 2 of 2
1 0 1 0 0 0 1 1 0 0 1 0 1 1 1 0
0 1 1 0 1 1 0 0 0 1 1 0 0 0 0 1

Course Introduction

Classification	<i>Information contained in this instruction is UNCLASSIFIED. However, certain methodologies are Law Enforcement Sensitive.</i>
Introduction	NITRO is a three-week course consisting of 14 days of lessons, incremental practical exams and a final practical exam.
Objective of this Course	NITRO is designed to introduce the officer to basic network intrusion investigation techniques.
Learning Outcomes	After completing this course the trained officer should be able to successfully conduct a network intrusion investigation.
Course Protocols	<p>Information contained in each section of this student book is presented in sequential order so that knowledge gained from later lessons is built on a foundation of what was learned earlier. Other course protocols include the following:</p> <ul style="list-style-type: none">• Practical exercises – Instructors will provide directives and handouts for practical exercises completed in the lab.• Appendices – Include course related materials provided by the instructors.
Practical Exercises	<p>Practical exercises in NITRO are hands-on. Each exercise is instructor-directed. In the exercises, students will:</p> <ul style="list-style-type: none">• Perform network wiring and connecting activities• Conduct manual log analysis• Use automated log analysis tools• Perform “Live” network gathering and analysis activities <p>In addition, every morning the students will have an opportunity to ask questions and/or review materials discussed on the previous day. This allows instruction to remain fresh and aids students with building practical connections to the training.</p>

This page intentionally left Blank

Network Intrusion Responder Program (NITRO)

Table of Contents – Book II

Topic	Page
Module 7 – Report Writing	7-1
Lesson 1 – Defining an Intrusion	7-3
The Forensic Report.....	7-4
Examiner Notes.....	7-5
Forensic Reporting.....	7-6
Title Page	7-8
Items Analyzed	7-10
Relevant Software.....	7-11
Glossary	7-12
Details of Findings.....	7-13
Items Provided	7-16
Creating a Hyperlink in Microsoft Word.....	7-17
Lesson 2 – Cyber Crime Interviews	7-19
Cyber Crime Interviews.....	7-20
Interview Process	7-22
Module 8 – Legal Issues.....	8-1
Lesson 1 – Search Warrants	8-3
Search Warrants.....	8-4
Search Warrant Exceptions.....	8-10
Consent Searches	8-11
Search Incident to Arrest or Apprehension.....	8-14
Other Search Warrant Exceptions.....	8-16
Lesson 2 – Internet Service Providers	8-19
Legal Framework	8-20
Express Consent.....	8-24

Topic	Page
Written Consent	8-26
Preservation Letters	8-29
Subpoena.....	8-30
Search Warrant.....	8-31
Available Data	8-32
 Module 9 – Fundamentals of Log Analysis.....	9-1
Lesson 1 – Understanding Network Traffic	9-3
Overview of Network Traffic	9-4
Investigation Techniques	9-5
Lesson 2 – The Scientific Method and Intrusion Analysis	9-9
Overview of the Scientific Method.....	9-10
Digital Forensic Analysis and the Scientific Method	9-12
Lesson 3 – Observing Intrusion-related Activity and Generating a Hypothesis .	9-15
Common Observations.....	9-16
Hypothesis Formation.....	9-19
Incident Classification	9-21
Lesson 4 – Predicting the Nature and Location of Intrusion Artifacts.....	9-25
Predicting the Nature and Location of Intrusion Artifacts.....	9-26
Relating Observed Events to Network Services and Traffic Types	9-27
Mapping Observed Activity to Traffic Flow	9-29
Using Traffic Flow and Service Type to Predict Artifact Location.....	9-33
Lesson 5 – Using Log Analysis to Evaluate an Intrusion Hypothesis	9-37
Hypothesis Evaluation	9-38
Acquiring Target Log Files.....	9-39
Reviewing Target Log Formats	9-40
Establishing Search/Extraction Criteria.....	9-41
Searching Target Logs and Extracting Relevant Data.....	9-42
Recording and Correlating Findings.....	9-43

Topic	Page
Keeping Track of New Leads	9-45
Module 10 – Log Sources	10-1
Lesson 1 – Windows Log Sources	10-3
Windows Logs	10-4
Windows Services Logs.....	10-6
Lesson 2 – Linux Log Sources	10-9
Linux Logs	10-10
Lesson 3 – Solaris Log Sources	10-13
Solaris Logs	10-14
Lesson 4 – Log Searching	10-15
Log Searching	10-16
Regular Expressions.....	10-17
Regular Expressions: Literal Characters.....	10-18
Lesson 5 – IDS Logs	10-19
IDS Logs	10-20
Module 11 – Log Analysis	11-1
Lesson 1 – Binary Traffic Analysis.....	11-3
Introduction to Wireshark	11-4
Converting Binary Logs to Text Format.....	11-5
Filtering and Searching in Wireshark	11-6
Filtering Data during Capture with Wireshark	11-7
Filtering Displayed Data in Wireshark	11-8
Colorizing Data Using Filters in Wireshark	11-14
Searching in Wireshark.....	11-16
Generating Statistics with Wireshark.....	11-17
Exporting Data from Wireshark.....	11-22
Lesson 2 – Manual Log Analysis	11-23
Filtering and Searching Text Logs.....	11-24
Deciding What to Search For.....	11-25

Topic	Page
Example Log	11-26
Lesson 3 – Automated Log Analysis Tools	11-29
What is Sawmill?	11-30
Installing Sawmill	14-31
Network Log Analysis Using Sawmill	14-38
Module 15 – LiveWire Investigations	12-1
Lesson 1 – Data Collection	12-3
Locating Physical Devices	12-4
Attaching Storage Equipment	12-6
Lesson 2 – Introduction to LiveWire	12-9
Live Digital Investigations.....	12-10
LiveWire Installation	12-13
LiveDiscover Installation.....	12-14
Updating LiveWire	12-16
Updating LiveDiscover	12-17
LiveWire Initial Setup.....	12-19
Lesson 3 – LiveDiscover	12-31
LiveDiscover Network Scanning.....	12-32
Lesson 4 –Volatile Data Analysis	12-39
LiveWire Initial Inquiry	12-40
System State.....	12-49
Current User Activity.....	12-55
Active Network State	12-68
Lesson 5 – Evidence Collection	12-73
File System Status.....	12-74
Physical vs. Logical	12-78
Collection and Preservation	12-84
Hashing	12-88
Lesson 6 – Malicious Code Analysis	12-93
Malicious Program Search.....	12-94

Topic	Page
Lesson 7 – Alternate Data Collection Tools	12-99
Windows Forensic Toolkit.....	12-100
Helix.....	12-103
Appendices	A-1
Appendix A – Intrusion Report Template	A-1
Appendix B – Volatile Data Collection	B-1
Appendix C – Understanding Computer Hardware	C-1
Appendix D – Data Storage Components	D-1
Appendix E – Input/Output Components	E-1

Module 7

Report Writing

Overview

Investigations require comprehensive reporting that documents actions and summarizes findings. The best reports are clear, concise, accurate, and report only information relevant to the facts of the case.

Purpose of this Module

The purpose of this module is to introduce you to an acceptable format and strategy for reporting. You will learn how to summarize the steps and findings of an investigation involving digital data.

Objectives

After successfully completing this module, you will be able to:

- Discuss the importance of writing an organized, clear, concise and accurate report
- Write an organized, clear, concise, and accurate report
- Discuss appropriate questions and interviewing techniques for computer crimes

In this Module

The following table lists the contents of this module.

Lesson	See Page
Lesson 1 – General Report Writing Techniques	7-3
Lesson 2 – Cyber Crime Interviews	7-19

This page intentionally left blank.

Lesson 1 – General Report Writing Techniques

Introduction Forensic reports involving the analysis of digital evidence should address the same basic information. No matter how well an investigator conducts analysis, it is of little value if results cannot be reported in an organized, clear, complete and concise manner.

Purpose of this Lesson The purpose of this lesson is to provide guidance for generating a document to report the forensic analysis results of digital evidence.

Objectives After completing this lesson, you will be able to create a forensic report that:

- Discusses the purpose and need for forensic analysis
- Explains what physical and/or logical evidence was analyzed
- Defines programs, terms, and their relevance
- Explains findings in an orderly manner
- Associates relevant evidence with users

In this Lesson The following table lists the contents of this lesson.

Topic	See Page
The Forensic Report	7-4
Examiner Notes	7-5
Forensic Reporting	7-6
Title Page	7-8
Items Analyzed	7-10
Relevant Software	7-11
Glossary	7-12
Details of Findings	7-13
Items Provided	7-16
Creating a Hyperlink in Microsoft Word	7-17

The Forensic Report

Introduction

The forensic report is the culmination of a process often involving intensive and painstaking work. It should reflect the time, effort and professionalism involved in building the case and acquiring the information. No matter how overwhelming conclusive the evidence is in determining guilt or innocence, if the evidence is not presented in an organized, clear and concise manner, it may be of little use to its intended audience.

Ultimately you should consider the report a reflection of yourself, your skills, attention to detail, experience and work ethic. In this respect, the report is much like a resume. It deserves the same attention that you would put into the development of your resume. The report should be well organized, include only relevant information, and be free of grammatical, punctuation and spelling errors. The recipient should be able to read it one time and have a very clear understanding of the message you are trying to convey. If the reader cannot do that your forensic report may be disregarded or discarded.

The consequences of such a fate for a resume are obvious but the impact is limited to the individual. The consequences of that happening to a forensic report could be far more reaching. If the evidence to convict a child predator was apparent but discarded for lack of organization and presentation the consequences could be felt by an innocent child or many children in the future.

Again, consider the report a reflection of your professionalism and develop it as such. It is a professional document and could very well be one of the most important reports you will ever write.

Examiner Notes

Importance of Documentation

The documentation that is created during the analysis process provides the basis for the examiner to report the results of the case. Remember, all written notes and documentation created during the investigation should be preserved and may be discoverable in court.

Examiner notes taken during the execution of a forensic examination along with the final report of your findings are the foundation on which many digital media-related cases are built.

Note Taking

Note taking is an essential part of a forensic examination. Your notes help create a document that will provide a record of all of the procedures and processes performed. The examination notes should present a clear timeline of the actions taken and the results of those actions. For example:

January 15th, 2008

0800 – Performed a signature check on the suspect media, Item #1, Partition1, NTFS, 6.0GB. There were a total of 243 file signature mismatches identified.

0835 – Reviewed signature mismatched files and found 23 files of interest ...

0900 – Used Adobe Photoshop 8.5 to review image files for existence of layered images. None were found.

Properly recorded notes provide a repeatable roadmap of your examination. Another examiner should be able follow your notes to reproduce the same results obtained in the original exam.

You should number, date, and initial all note pages using the [page #] of [total # of pages] numbering schema to account for all note pages. It is not uncommon for extensive periods of time to pass between the time of the examination and prosecutorial action.

Thorough note taking will help ensure that you can accurately testify to actions taken during the examination. When taking examiner notes, always follow the rule, “If it wasn’t documented, it didn’t happen.”

Forensic Reporting

Introduction

A forensic analysis report should contain all of the relevant evidence that you find during your examination. Your final report must clearly identify persons related to the examination including you, the requestor, suspects, and any other pertinent individuals.

Your report should:

- Provide details about the purpose for the forensic analysis
- Describe the physical and/or logical evidence analyzed
- Define related programs, terms and their relevance

Most importantly, the forensic report must clearly and concisely explain the items of evidentiary value that were found on the suspect media as a result of your analysis. It must also identify the location and relevance of the items of evidentiary value as relating to the reason for the analysis and/or the investigation.

Report Contents

Each completed forensic report should always contain the following information, although the headings used within your report may vary from organization to organization:

- Report heading
- Support requested, reason or purpose for analysis
- Summary of findings
- Digital media analyzed
- Analysis/Suspect Software Listings
- Glossary of Technical Terms
- Detail of Findings
- Items Provided

Your forensic report should define all technical terms using common language that non-technical investigators and prosecutors can understand. It should clearly indicate relevant items you discovered, how they were discovered, where they were located, and how the evidence relates to the case and/or investigation.

Forensic Reporting, continued

Note Taking versus Reporting

It is important to distinguish between note taking and reporting. The forensic report is the final result of your analysis and its focus is to respond to the allegations or charges that led to the investigation, defined by:

- The request or purpose for the analysis
- Information compiled as the case progressed including search authorization documents
- Information provided by other sources, legal, victims, and informants
- Your experience as a forensic examiner

The report should include those items that directly relate to and are relevant to the allegations or charges in the request for analysis.

Notes, however, should include references to all steps taken by the analyst during the examination with either positive or negative results. The primary goal of note taking is to document all steps taken as well as to serve as a chronicle of the examination. It may be used as a reference later during prosecution.

Title Page

Contents of the Title Page

The title page provides an overview of the case, including:

- Report Header
- Support Requested
- Current Case Status
- Summary of Findings

Description

In the table below, you will find standard information that should be included in the title page of all reports. A sample report follows this chart.

Item	Description
Title (To:)	Indicates report's recipient and date. This information is usually directly related to the charging document or request for service.
From	Identifies report's author, including name, organization, and contact information.
Subject	Indicates the type of work performed, as well as any suspect and/or local case reference information.
Support Requested or Purpose for Analysis	Concisely states what charges or allegations were addressed by the analysis, charging jurisdiction or entity, with regard to a specific investigation.
Status	Indicates the current case status, usually Open or Closed. This may change the nature from a report to a status update.
Summary of Findings	A short narrative describing the type and nature of any evidentiary items located during analysis with respect to the specific allegations or charges. The failure to locate items that support the allegations should also be indicated here, as well as items that may exonerate.
Footer	Includes author's name, title/rank, and a "Released" field for approval signature. Footer may include a legal disclaimer. All pages of the report should be numbered in the Page X of X format.

Title Page, continued

Title Page Example

Computer Forensic Analysis Report	
MEMORANDUM FOR Trooper William Barksdale, MD State Police	January 30, 2008
FROM:	Trooper Michael Ghaler, Forensic Examiner MD State Police Forensics Laboratory Pikesville, Maryland 21208-3899
SUBJECT:	Forensic Media Analysis Report Subject(s): John M. Smith Born: October 11, 1960 Case Number: 2008-0123A
1. Support Requested	
Trooper Barksdale, MD State Police requested the examination of the submitted digital media seized as the result of an ongoing criminal investigation related to the theft and subsequent sale of weapons from a U.S. Army facility in Baltimore, Maryland. The request asked that the submitted media be examined for any evidence related to the theft and/or sale of explosives, particularly C4, and the existence of possible co-conspirators who participated in, or facilitated the unlawful activity.	
2. Status: Pending (or closed)	
3. Summary of Findings	
The examination of submitted media, which was conducted during the period of 2008/01/14 and 2008/01/18 resulted in the discovery of numerous image files depicting the type of explosives indicated in the Request for Analysis. A number of the image files also showed the suspect, SMITH, in possession of the materials. Additionally, a review of Internet activity resulted in the discovery of a number of instances where the suspect's computer was used to access Internet auction sites in an attempt to sell explosive materials. While the suspect received numerous email correspondences from individuals apparently responding to his auction advertisements and participated in newsgroup conversations regarding the sale of explosives, little information was found which indicates the existence of a conspiracy between the suspect and any other individuals in this activity.	
Trooper Michael Ghaler Computer Forensic Analyst	Released by: _____

Items Analyzed

Introduction The Items Analyzed section of the report describes in detail the analyzed physical and/or logical evidence. It should always include the original *and* verified hash values of all evidence items.

Physical Items The Items Analyzed section of the report may contain a reference to the actual physical evidence. Give a detailed description, including:

- Manufacturer
- Model, serial, and part number (when possible)
- Item description
- Any specific markings

Logical Items The Items Analyzed section of the report details images sent to you for analysis as well as any “original” evidence items you may receive for examination. In addition to listing the physical containers, you should also list the image files. Do this by original file name and include any hash or other validation mechanism. An example follows.

Items Analyzed Example

Items Analyzed

Tag # 01 Western Digital Caviar 31600 Hard Drive
Serial #: WT2891586134
Size: 40GB
Hash Provided: 1234567890ABCDEF1234567890ABCDEF

Tag # 02 One Memorex DVD-R disk containing image files derived from a Fujitsu M1636TAU Hard Drive, Serial #: 08613105, Size: 1226MB:

Suspect_dvd.e01
Suspect_dvd.e02
Suspect_dvd.e03
Suspect_dvd.e04

Hash Provided: ABCDEF1234567890ABCDEF1234567890
Verified Hash: ABCDEF1234567890ABCDEF1234567890
Hash values indicated above are for the entire device image. Individual file hash values were not provided.

Relevant Software

Introduction

This section of the report identifies the software found on the evidence media that is relevant to the case as well as the identity of forensic software used to perform the examination.

Analysis Software

List all software applications used during the forensic examination to process or analyze the suspect media, including the primary analysis tool. Be sure that all applications in this listing include the appropriate software version information and a brief description of the software's functionality or use.

Suspect Software

Identifying the software on a suspect's machine is a necessary component of the case. Also include any software that may have created and/or interacted with data of evidentiary value located during the forensic examination. Include any software that you can identify as being deleted from the suspect's machine. Pay particular attention to software which is commonly used to hide data or securely erase data from the device. Each software listing should include the:

- Software name and version
- The full path to where the application was located on the suspect media
- A brief description of the program functionality and how it relates to the Request for Analysis and/or investigation
- You should be prepared to further explain items in this listing during prosecution

Relevant Software Example

Analysis Software:

<u>Program Name/Version</u>	<u>Program Description</u>
LiveWire Investigator 2007	Live Forensic Analysis Suite
Sawmill Professional 7.2.14	Forensic Log Analysis Tool

Suspect Software:

<u>Program Name/Version</u>	<u>Directory Location</u>	<u>Program Description</u>
AOL Instant Messenger	C:\Program Files\AIM	Internet Chat Application
Microsoft Internet Explorer	C:\Program Files\Internet Explorer	Internet Browsing Application
Microsoft Outlook Express	C:\Program Files\Outlook Express	Email Application

Glossary

Glossary The Glossary defines any technical terms, document formats, and procedure details referenced within your report that may not be readily understood by the average non-technical reader. Only define the terms that are integral to the understanding of your examination findings as presented in the report.

Glossary Example

Glossary

Term	Definition
HTML (Hypertext Markup Language)	One of the authoring languages used to create documents on the World Wide Web.
IRC (Internet Relay Chat)	A multi-user chat system, which is real-time communication between two users via computer. Once a chat has been initiated, either user can enter text by typing on the keyboard and the entered text will appear on the other user's monitor. It is often used on the Internet.
JPG (Joint Photographic Experts Group)	A graphic image file format.

Details of Findings

Introduction

The Details of Findings section provides detailed information about any items of evidentiary value found on the suspect media during the forensic examination. The information in this section should be thorough, yet concise, and only contain details relevant to the request for analysis and/or the investigation. It should *not* contain information about processes executed that did not produce relevant information, unless the negative result is relevant. Information about these non-productive processes should be included in your notes.

Organization

There are several different ways to organize your report. You may find yourself using different organizational strategies on different investigations. You might, for example, organize a report on a case that contained many pieces of media by listing all of the items found on each piece of media as one section.

In other cases it may be more effective to organize the data by date and time. This approach works particularly well in cases where the organization of data such as e-mail, chat and downloads is easier to understand if organized in chronological order.

Another method of organizing your report would be to organize evidence by its relationship to a particular criminal charge and subject. This approach works well for criminal prosecution. It allows the prosecutor to quickly see the evidence, which is relevant to a particular charge and subject.

As shown in the following example, the evidence can also be ordered or segregated by device. This would be a good organizational choice for a case that has many pieces of media such as a large quantity of CDs or DVDs. This organizational strategy is not normally the best choice when all of the evidence is located on one device.

If the evidence has multiple partitions, you should further subdivide your results by individual partition. Make sure you detail the partition's file format and size. To further clarify evidence, you may also want to divide your findings by each user account or profile. You should detail important files, structure, data, and discrepancies. Explain your techniques, methodology, and the relevance of information in brief narrative statements whenever possible.

Details of Findings, continued

Organization, continued

If the analysis of a particular item of evidence did not result in the discovery of any items of evidentiary value, a simple statement should be included stating the negative results to ensure that the reader does not misinterpret an omission as the failure to analyze evidence.

Include in the report the techniques you used to locate or extract evidence. Later, you can refer to your report if asked in court how you found a specific piece of evidence. For example, if you found the file by manually inspecting the drive, then state this. Remember, not every search has to be done with an automated tool.

You should take time to develop an organized structure for your report before writing it.

Hyperlinks

If your report is included on a CD or other large media, use hyperlinks where possible to illustrate items of evidentiary value. Hyperlinks allow the reader to click on a description of the document, or data, when viewing the electronic copy of the report and display the actual file on the screen.

Details of Findings, continued

Details of Findings Example A skeletal example of the layout for the Details of Findings section follows.

Details of Findings

A complete analysis of all computer media analyzed (listed above) revealed the following information pertinent to the Request for Analysis and/or of evidential value to this investigation.

Analysis of the Tag #1, image files Suspect.E01 through E04 from 3.0262GB Maxtor hard drive revealed:

Partition 0, NTFS 2.0GB

[Detail your findings in this section relative to the indicated partition located on the hard drive specified. Use narratives, tables, lists, etc. to describe what you found, where, and relevance to the case, etc.]

Partition 1, FAT32, 1.0GB

[Detail your findings in this section relative to the indicated partition located on the hard drive specified. Use narratives, tables, lists, etc. to describe what you found, where, and relevance to the case, etc.]

Disk Free Space, 0.0262GB

[Detail your findings in this section relative to the indicated disk area located on the hard drive specified. Use narratives, tables, lists, etc. to describe what you found, where, and relevance to the case, etc.]

Analysis of Tag #2, black Memorex floppy disk revealed:

[Detail your findings in this section relative to any items of evidentiary value located on the floppy disk. Use narrative, tables, lists, etc. to describe what you found, where, and relevance to the case, etc.]

Analysis of Tag #3, SONY CD-R media labeled "My Plan" revealed:

[Detail your findings in this section relative to information located on the CD. Use narrative, tables, lists, etc. to describe what you found, where, and relevance to the case, etc. In the event of multiple session CDs, it may be necessary to further subdivide this information by specific session.]

Items Provided

Items Provided

This section details *all* of the physical items returned to the requestor with this report. It should include all of the items specified in the Items Analyzed section. Be sure to include items that were generated as a result of your analysis, such as a ZIP disk, floppy disk, CD-Rs, and hard copy documents.

Example

Items Provided

- 1) Tag # 01 Western Digital Caviar 31600 Hard Drive
Serial #: WT2891586134
Size: 40GB
- 2) Tag # 02 One Memorex DVD-R disk containing the image files listed below
derived from a Fujitsu M1636TAU Hard Drive, Serial #: 08613105, Size: 1226MB:

Suspect dvd.e01
Suspect dvd.e02
Suspect dvd.e03
Suspect dvd.e04
- 3) One CD Labeled "Findings of Case Number: 2004-0123A", which contains
information referred to in this document.
- 4) One printed document consisting of three pages, entitled "C4 – What's it good
for?" that explains the uses of C4 and methods to surreptitiously sell it using
publicly available Internet auction web sites.

Creating a Hyperlink in Microsoft Word

Introduction

Hyperlinking is a good way to direct readers to items referred to in a written report on electronic media. It is often not practical to display certain items in the actual report. For example, the report may reference an offensive, graphic image of child pornography. Some readers may not have the need to view the image. Others may be required to do so in order to confirm its existence. In this scenario, a simple hyperlink would allow those with the need to inspect the file while sparing others from having to view it.

Additionally, there may be so many items of interest that it would be too cumbersome to include all of the actual images in the written report. In this situation, a hyperlink could lead to an entire directory that contains multiple files.

These are both excellent examples of when hyperlinking would compliment your report.

Procedure: Create a Hyperlink in Microsoft Word

Use the following procedure to create a hyperlink in Microsoft Word.

Step	Action
1	Use the mouse to highlight the reference item.
3	From Microsoft Word's pull-down menu, select Insert > Hyperlink.
4	In the Insert Hyperlink dialog box, traverse the directory structure until you locate the file or directory to which you want to establish a link.
5	Click OK . The reference item should change color to indicate that it is now a hyperlink.

This page intentionally left blank.

Lesson 2 – Cyber Crime Interviews

Introduction

Interviews are an essential element of developing information that is relevant to a criminal investigation. When conducting a cyber crime investigation, investigators must prepare for the interview, develop rapport with interview subjects, ask questions that generate corroborative information and leads, and terminate the interview in a way that leaves the door open for further questions.

Purpose of this Lesson

In this lesson, you will learn guidelines that will assist you in conducting interviews for cyber crime investigations.

Objectives

After completing this lesson, you will be able to:

- Develop a plan to conduct interviews in a cyber investigation
- Explain the psychology and culture of the technology world and ways to apply that knowledge to the interview process
- Ask questions that will provide you with information that will assist the investigation

In this Lesson

The following table shows the contents of this lesson.

Topic	See Page
Cyber Crime Interviews	7-20
Interview Process	7-22

Cyber Crime Interviews

Overview of Cyber Crime Interviews

When most investigators think about cyber crime investigations, they think of the data stored on the digital media, volatile data from the physical devices associated with the incident, and logs from witness devices. All of these items contain information that is valuable to the case and should not be overlooked by an investigator.

However, none of this data completes the picture by itself because machines do not act on their own free will. In order to complete the puzzle that the investigator is trying to solve, he/she must obtain information from all of the people who are involved with the incident.

The witnesses and victims of a crime can help direct the investigation and assist the investigator in understanding the nature of the crime, locating more evidence, and identifying suspects. A proper interview of a suspect may assist in revealing the true scope of the investigation and provide the information needed to ensure the conviction of a suspect.

Interviews are an integral part of any investigation because the victims, witnesses, and perpetrators of a crime all have pieces of the puzzle that the investigator is trying to put back together. When confronted with an incident that involves complex technical issues and multiple sources of evidence, an investigator must skillfully navigate the human landscape to develop leads, confirm events, and obtain a complete picture of the crime.

Accusatory versus Non-Accusatory Interviews

There is a difference between an interview and an interrogation. The purpose of an interview is to gather information that will confirm events, develop suspects and leads, and identify facts that lead an investigator to identify the root causes of the incident being investigated.

Interviews are not focused on getting an individual to confess, but on developing information. In order to obtain the information he needs, an investigator develops rapport with the subjects and conducts interviews in a non-confrontational manner.

Interrogations focus on presenting the facts of the case to the suspect and eliciting statements that confirm the suspect's involvement in the incident.

Cyber Crime Interviews, continued

Accusatory versus Non-Accusatory, continued

Understanding the difference between an interview and an interrogation is critical to the investigator's success. If the investigator takes an aggressive and accusatory approach to communicating with interview subjects, there is a strong possibility that this type of interaction will close the door for any further effective communication.

The investigator must also take into account the perceptions of the interviewees. Many of the people that you will encounter during the initial interviews of a cyber crime investigation have little to no exposure to law enforcement outside of what they have watched on television. The subjects may also view any interaction with government agents through the lens of their personal attitudes, experiences, and beliefs about law enforcement.

In the hacker culture, law enforcement is often cast in a negative light and government agents are seen as dimwitted and overbearing. By taking these issues into account, an agent can prepare for the interview process and maximize the amount of information he can obtain.

Interview Process

Interview Process Interviews are often open-ended, free flowing conversations that develop information and enhance an investigator's understanding of the case. Following a structured process that progresses from basic to detailed questions provides more relevant results to the case agent.

Interviews during a cyber crime investigation are the first and best opportunity to determine what has happened and obtain clarification. The crime scene and evidence are not visible to an investigator during a computer incident. Therefore, investigators need to gather information about the victim network and machines from other people. The individuals who control the machines understand the geography of the digital environment. An investigator must obtain this information from these people.

The interview process has been broken down into the following model which can be used to structure the approach taken when interviewing subjects from a technical background.

- Planning/Research
- Opening/Rapport
- General Questioning
- Detailed Questioning
- Interview Termination

Planning and Research

Prior to conducting an interview, it is a good idea to conduct preliminary research on the interviewee and the organization that he or she represents. Understanding a person's training, experience, area of expertise, the nature of his or her daily business, and role within an organization can help the investigator draft appropriate questions and an interviewing approach that will be successful in a specific corporate culture.

Investigators who are knowledgeable and comprehend a technical subject's area of expertise present a professional and approachable image to the interview subject. The knowledge gained through planning and research also sets the stage for establishing a situation where the interviewer can ask more intelligent questions.

Interview Process, continued

Opening/Rapport The initial contact with subjects in a cyber crime investigation is critical to obtaining cooperation during an interview. Showing respect for the needs of the individual and the organization and explaining the importance of the interview questions are vital steps in obtaining a subject's cooperation. You should explain to the interview subjects that you will need their assistance in determining the facts of the case and understanding some of the issues that are involved.

General Questioning Start with open-ended questions that allow the subject to relate to you his or her knowledge of the events. These questions should develop a general framework of the incident that is derived from the subject's personal knowledge of factual events. During a technical investigation, it is important that an investigator obtain the general outline of the events that transpired in order to document information that is relevant to the case.

You should obtain the following information:

- Key incidents that brought the situation to light
- Hardware that was involved (routers, firewalls, IDS)
- Specific individuals who were involved in the incident and the physical actions they took with any evidence
- Physical locations of effected machines and people
- Technology that will need clarification during the Detailed Questioning phase

Interview Process, continued

Detailed Questioning

After obtaining a general picture of the incident, the investigator can probe deeper into issues that are relevant to the case. The investigator can ask the subjects to describe in more detail how they became aware of the incident, what physical actions they took, and how they came to any conclusions. These questions are important to obtaining a complete picture for the investigator.

The following information should be pursued and collected at this stage:

- Software and hardware model numbers and versions
- Network monitoring and logging setup
- Collection of current logs
- Network diagrams
- User policies
- List of people who had logical and physical access
- Ownership and authorized access of systems
- User account information
- Statements of how the systems are used
- Individual access to relevant systems and/or data
- Specific commands or tools used during the discovery of the incident
- Security video or proximity card logs

Interview Process, continued

Interview Termination

It is not uncommon for the investigator to go back for more details as his knowledge of the incident grows or if he needs assistance explaining a complex technical subject to other people. The investigator should ask the subject for contact information in case the investigator needs to clarify or elaborate upon material discussed during the interview. By terminating the interview in this manner, the investigator has another opportunity to speak with people that may become suspects at a later time.

Interview Psychology

Investigators should take into account the subject's perspective and culture. Many times people in the computer field are distrustful of government agents and often have a skewed perception of what sort of evidence an investigator may need or want. System administrators from a college campus or library will typically have very different values and perceptions than a system administrator that has worked in a DoD environment. By identifying the work place environment, professional experience and culture, and common characteristics that influence the interviewees, the case agent can determine the best approach to building rapport with the individuals that hold insight to key elements of the case.

Establishing the right tone for an interview is important in making sure that the interviewer is obtaining all of the necessary information that is available from the subject.

The approach an investigator takes to a case differs depending on whether the investigator initiates contact with an organization for information or whether the investigator is called in by the organization who is a victim of the crime. This issue will determine how the investigator structures his approach to the interviews.

When an organization calls in an investigator, most of the organization's key players are already involved in the incident. Upper management will typically support the initial stages of the investigation. In this situation, you should obtain written documentation from everyone involved in the incident and begin working on controlling the investigation. It is important to minimize the mishandling or destruction of evidence and the spread of information about the investigation to keep potential suspects in the dark.

Interview Process, continued

Investigator Initiated Contact

When an investigator is tracking down a cyber incident, the case will often lead to an external organization that may have computer logs or actual systems that are relevant to the case.

First the investigator must identify a point of contact within that organization. This can be done by looking up registration information for the company via an Internet search engine and then identifying the legal counsel and network administrator. There are also a wide variety of state, federal, and private registries that contain this information and will provide it to law enforcement agencies. You can also find several Internet service provider lists that include contact information for major communications providers such as AOL, Microsoft, Google, Yahoo, and others.

Then investigators must decide how they will contact a complex organization that they may know nothing about. These are some of the considerations that an investigator should take into account:

- Is the system administrator or someone with root access responsible for the activity?
- Do they know the suspect and are they relaying information about the investigation to their associate?
- Are there any regulatory or legal barriers to the organization giving information to the investigator (ECPA, PPA, FERPA, HIPPA, Organizational Policy, etc.)
- How critical is the evidence held by the company to the case? Is a life or national security at stake?
- Does the investigator have access to request the legal instruments such as subpoenas, search warrants, FISA, or other items in a timely manner?
- Does the investigator have the technical knowledge to handle the systems or logs that will be obtained from the scene or will he need additional assistance?
- Will the evidence be contained all at one physical location or spread out over a national or global network infrastructure?
- Does the investigating agency have the manpower, technical knowledge, and resources to obtain what they need from the target organization?

Interview Process, continued

Investigator Initiated Contact, continued

When an investigator reaches out to an organization, it is best to identify the senior management, legal counsel, and technical heads prior to making contact. If the investigation allows, you should start at the top of the organization's hierarchy and work down towards individuals who are responsible for managing or using various computer and network services. An investigator can address legal issues, business impact, and concern for the organization's needs in order to gain more cooperation.

If an investigation could be compromised by reaching out to the organization in this manner, the investigator should avoid "tipping off" the organization or individuals. The investigator can then do the research and case preparation to obtain the most effective legal authority that will allow him to gather evidence and interviews in the most effective and efficient manner possible.

Organization Initiated Contact

If the victim organization is reaching out to the investigator to report an incident, the investigator will more than likely have cooperation. At this point, the victim performed an internal investigation that may or may not have been done properly. The investigator will still need to follow up to ensure the accuracy of any information provided by the victim organization. Again, properly document any actions taken by people at the company.

It can be very important to work with the organization's legal counsel and network administrators. The legal counsel can assist with collecting information that is governed by organizational policies.

Interview Process, continued

Organization Initiated Contact, continued

Many organizations are ignorant of the investigation's process and do not use proper evidence handling procedures. Common practices used by organizations to minimize or repair damage caused by an incident will often hamper the investigation. Some of these issues are:

- When a machine has been compromised it is common practice to just restore the machine from a base image.
- Many times untrained system administrators or management will work on live systems and alter system artifacts and time lines.
- There is more concern for system/network integrity compared to retaining any evidence or identifying the cause of the incident.
- Chain of custody issues are often not followed
- Lack of network maps and diagrams
- Incomplete knowledge of system functions and passwords
- Poor documentation of system builds, organizational policies, or security controls

Investigators should address these issues in advance to minimize problems with this type of situation.

Witness and Victims

These individuals often provide the investigator with the initial foundation of a case and the information needed to build the framework for the investigation. A system administrator or a user of an online database may be the first to notice strange system behavior or altered data in a critical database. A relative or friend may discover suspected child pornography images on a computer and report it to the police. In either situation, the investigator will need to elicit information from people to obtain facts that will lead them to a suspect.

Interview Process, continued

Witness and

Victims, continued

Investigators must keep in mind that anyone could be a suspect. During the initial response and interviews, an investigator needs to obtain clear and concise written documentation from the people involved with the incident. This documentation should consist of any actions taken by the witnesses and events observed by them during the incident. Your documentation is an important part of freezing the crime scene and creating a permanent record of events that occurred during the discovery of the crime.

Investigators should keep in mind that the people who are being interviewed may have had little contact with law enforcement and may not recognize the type of evidence needed to further the investigation. It is the investigator's job to develop a rapport with the interview subjects and guide the interview in a direction that will obtain the most relevant and complete information possible.

Issues to be Addressed During Interviews

There are various issues that should be addressed during interviews with witnesses. Here are some important questions:

- Are system administrators or people with administrative powers potential suspects? If so, the investigator must move quickly and efficiently to build the case. If immediate action is not an option, the investigator must ensure that the suspect's access to evidence or sensitive material is removed. Covert tactics may become necessary if you have the proper administrative approval, legal documents, and technical capabilities in place.
- Obtaining witness statements about the suspect's access to the machine or data is also important. Investigators should supplement any network logs and/or forensic analysis with as much traditional detective work as possible.
- Do people share machines, use passwords, or share passwords?
- Have there been any strange phone calls or repairmen asking for user accounts, passwords, or other sensitive technical information? If so, these incidents may have been a suspect trying to obtain information about the victim organization in order to commit the crime.

Interview Process, continued

Suspects

Traditionally, investigators interview a victim and eventually a suspect. However, in cyber crime cases, it can be unclear who is a victim or who is a suspect. The investigator must use analytical skills to determine who should be interviewed during an investigation.

Module 8

Legal Issues

Overview No matter how incriminating the evidence may be, all computer crime investigations must adhere to established legal principles. If legal standards are not met, the case could be jeopardized and even dismissed, thus allowing a perpetrator to walk free.

Purpose of this Module The purpose of this module is to familiarize you with some basic legal issues that must be considered when conducting an investigation involving digital data.

Objectives After successfully completing this module, you will be able to:

- Explain some of the legal issues involved in a digital investigation
- Employ practices during an investigation that that will pass legal challenge

In this Module The following table shows the contents of this module.

Lesson	See Page
Lesson 1 – Search Warrants	8-3
Lesson 2 – Internet Service Providers	8-19

This page intentionally left blank.

Lesson 1 – Search Warrants

Introduction

The search of a person or a location requires either a search warrant or a valid exception under the 4th Amendment to the U.S. Constitution.

Purpose of this Lesson

The failure to comply with the provisions of the 4th Amendment to the U.S. Constitution may result in the exclusion of valuable evidence at trial because the evidence was not legally seized. Investigators must document their authority to search and seize and be prepared to effectively articulate the probable cause that justified the search. In this lesson, you will learn about search authorities and how they are obtained.

Objectives

After completing this lesson, you will be able to:

- Explain how the 4th amendment of the United States Constitution is applied by the government
- Recognize situations in which the investigators may search or seize without a warrant
- Discuss the types of consent and their requirements

In this Lesson

The following table shows the contents of this lesson.

Topic	See Page
Search Warrants	8-4
Search Warrant Exceptions	8-10
Consent Searches	8-11
Search Incident to Arrest or Apprehension	8-14
Other Search Warrant Exceptions	8-16

Search Warrants

Introduction

The 4th Amendment to the United States Constitution reads:

“The right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures, shall not be violated, and no Warrants shall issue, but upon probable cause, supported by Oath or affirmation, and particularly describing the place to be searched, and the persons or things to be seized.”

This amendment sets forth the foundation upon which all search warrants are justified and the standard by which the legality of a search warrant is judged.

Relevance

Search warrants are a common tool used by prosecutors and investigators in any criminal investigation. They establish the authorization for the search and seizure of evidence with the court or designated approval authority prior to any search or seizure.

Search warrants provide the most reliable means of obtaining evidence in an investigation. Although search warrants are subject to legal challenge, if properly crafted and executed, they are difficult to overcome. The use of a Search Warrant is preferred by the U.S. Supreme Court whenever investigators have probable cause to believe a crime has been committed and are seeking search authority.

Obtaining a Search Warrant

Investigators seeking to obtain search warrants should become familiar with Federal Rules of Criminal Procedure Rule 41, “Search and Seizure.”

An investigator who requests a search warrant must establish by sworn affidavit the following key pieces of information:

- Description of the place to be searched
- Concise description of the item(s) being sought
- Probable Cause or facts that support the belief that the items being sought are located in the place described

Failure to establish any of these facts can result in the search warrant being overturned or not issued in the first place.

Search Warrants, continued

Obtaining a Search Warrant, continued

A current federal search warrant template is found as an editable PDF format at:

<http://www.uscourts.gov/forms/AO093.pdf>

The warrant is typically accompanied by an attached affidavit. The affidavit will normally contain the following components:

- The affiant's statement of probable cause
- Attachment A – Place To Be Searched
- Attachment B – Items To Be Seized

There is no formally required format for any of these documents, although various agencies and jurisdictions may follow formats that have been developed over time.

The warrant itself generally provides brief statements of these three components and refers to the affidavit for greater detail. For example, the warrant contains the text “In the Matter of the Search of” in the upper left. If the location to be searched is a residence, the warrant might read:

In the Matter of the Search of:

123 Patriot St, Foxboro, MA 02345, further described in Attachment A

The United States Department of Justice, Computer Crime and Intellectual Property Section (CCIPS), publishes a very useful guide called *Searching and Seizing Computers and Obtaining Electronic Evidence in Criminal Investigations*. This publication is available in hardcopy, and online at:

<http://www.cybercrime.gov/s&smanual2002.htm>

Appendix F, “Sample Language for Search Warrants and Accompanying Affidavits to Search and Seize Computers,” of this publication is the de facto standard for warrant and affidavit language in federal cyber investigations.

Search Warrants, continued

Description of the Place to be Searched

The 4th Amendment requires the investigator to concisely identify and define the search location in physical terms. This is often difficult to do in a cyber crime investigation as the physical and virtual worlds may not share the same physical space.

The investigator should be as careful as possible when describing the place to be searched. The goal is to define the boundaries of the search location in such a manner that the area within those boundaries may be searched, but the search is not overly broad.

The following would be considered an “overly broad” description of a place to be searched:

All property owned by John Smith.

A better description may read:

The residence located at 2021 Colony Drive, Podunk, AFB, MD described as a two-story wood-frame single family residence located on the south-east corner of the intersection of Colony Drive and Athens Way. The residence is painted off-white with brown trim and is distinguished by a brick mailbox in front with the numbers “2120” on the front of the mailbox facing Colony Drive. The location includes an attached single-car garage and detached storage shed located behind the residence inside a fenced back yard.

Search Warrants, continued

Description of the Item(s) to be Seized

Investigators crafting a search warrant affidavit should carefully consider the description of the items they wish to seize. In most cyber crime cases, the item sought is basically information.

If investigators can describe the information they seek through the search, they can more easily articulate the different forms the information may take and the different storage media upon which that information may reside. This type of approach helps clarify the description of the items to be seized. A description of items to seize in a search warrant for a intrusion case might resemble the following:

...for the seizure of exploit tools, account information, passwords related to XYZ Corporation. This information may be stored in physical documents, notes, papers, electronic storage media including but not limited to: computer memory, hard disk drives, Flash memory cards, floppy diskettes, smart cards, memory stick, secure digital media or other removable electronic storage media, cellular phone storage devices and personal digital assistant devices, compact disks, DVD disks and similar optical storage media as well as indicia of ownership.

The investigator must establish probable cause in a sworn affidavit for each item to be seized in the search warrant.

Search Warrants, continued

Probable Cause

The 4th Amendment states “...and no Warrants shall issue, but upon probable cause, supported by Oath or affirmation” This means that there must be a specific reason why each item is related to the investigation and that reason must be addressed in the investigator’s affidavit in support of a search warrant.

The investigator can base his reason for seeking an item on a variety of facts as well as personal experience and training. In the affidavit, the investigator should explain his experience and list all related training regarding his expertise in the field. Once expertise is described, the investigator can articulate the probable cause for the seizure of items listed in the search warrant.

Here is an example of a statement in an affidavit:

Based upon my experience and training in investigation of computer network intrusion cases, I know that computer network intrusion suspects typically keep notes and background information related to victim systems and the compromises of computer systems and networks. They also typically use automated software called “exploit tools” on computers to attack computer systems and networks. Exploit tools often capture or return data concerning the victim system and that data can be stored on a variety of storage media accessible to the intruder. For this reason, I believe that the location will likely contain storage media as well as computer and telecommunications equipment necessary to exploit computer systems. Additionally, many computer and digital devices are purchased through commercial sources and records of those purchases are often kept by the owner for warranty and accounting purposes indicating ownership of the device(s)...

Search Warrants, continued

Execution of a Search Warrant

Once a search warrant is signed by a judge, you can serve the warrant, initiate the search, and collect evidence described in the warrant. Most court jurisdictions require that the investigator submit a Search Warrant Return document to the court describing the items actually seized at the scene.

The search warrant gives the investigator the lawful authority to enter the described premises and conduct the search irrespective of the objections of the owner or occupant. It is not unusual for the owner's attorney to show up during the execution of the warrant and demand the search be stopped. If this happens, you should politely refer the attorney to your prosecutor for discussion and continue the search. The owner or his attorney does not have the right to obstruct or interfere with your execution of the warrant.

Surreptitious Execution of a Search Warrant

In some cases, such as those involving organized crime figures, or violent individuals, you may want to execute the warrant without the knowledge of the individual and without public disclosure of the warrant or affidavit until trial. In these cases, a surreptitious entry warrant may be requested and issued.

The surreptitious entry warrant authorizes the investigator to enter the premises and conduct the search without the individual's knowledge. Affidavits, search warrants, and the Search Warrant Return are kept under seal and not made public.

A similar physical search authority can be issued under the FISA provision in accordance with the USA PATRIOT Act.

Search Warrant Exceptions

Introduction

Through several rulings, the U.S. Supreme Court has interpreted specific exceptions for the 4th Amendment requirement to obtain a search warrant.

Relevance

Investigators need to understand the circumstances under which they are authorized to search an individual or premises without a warrant.

**Warrant
Exceptions**

Within certain limitations, an investigator may search an individual or premises without a warrant in the following circumstances:

- Consent
- Stop and Frisk
- Search Incident to Arrest
- Immediate threat to life or serious bodily injury
- Immediate threat of the destruction of evidence
- Fresh pursuit
- Plain view
- Vehicle searches
- Custodial searches
- Border searches

Of these, consent searches are the most common exceptions to the warrant requirement.

Consent Searches

Introduction	A person may waive his rights under the 4 th Amendment and consent to the search of his person or items under his control.
Relevance	Consent is a powerful tool for obtaining search authority. Evidence obtained during a consent search is admissible in court as long as the investigator obtained the proper consent. Investigators should understand how consent is granted and the limitations on consent searches.
Owner Consent	<p>A property owner has the legal authority to authorize the search of the premises as long as certain requirements are met:</p> <ul style="list-style-type: none">• Consent must be voluntary and not coerced• Consent must be informed• Consent can be withdrawn at any time• Consent can be limited <p>When a consent search is granted, the consent should be obtained in writing, signed by the consenting party, and dated with a known good local time and location. Most agencies have a consent search form for this purpose.</p>

Consent Searches, continued

Voluntary Consent For a consent search to be valid, the owner must consent freely and knowingly. The investigator may ask for consent to search. If this permission is granted by the owner, the investigator may search the property or premises legally. Some jurisdictions require that the request for consent to search be predicated on definable suspicion.

The investigator cannot coerce the owner by threats, intimidation, or power of authority into consenting to a search.

Informed Consent Not everyone is capable of giving consent to search. Some are legally and mentally incompetent and unfit to make such a decision. Those with clinically diagnosed mental conditions and severe health problems that affect their judgment are unable to intelligently consent to a search of person or property.

Minors are also generally not trusted to give consent. Many jurisdictions debate whether a juvenile can give consent without parental approval. As an investigator, you should be familiar with the court decisions in your area concerning informed consent to search.

**Withdrawing
Consent**

A person who waives his rights under the 4th Amendment can reassert those rights at any time by simply telling the investigator to stop the search. Once told to stop, the investigator must stop the search unless he possesses some other legal authority to continue the search.

Evidence located before the consent to search is withdrawn is admissible and can be retained by the investigator for further analysis. For example, if an investigator made a forensic copy of evidence during the consent search, he has the right to retain the forensic copy and examine it even though consent is later revoked.

For these reasons, the investigator should prioritize the places to search and items to seize when operating under a consent search. This helps to maximize the effectiveness of the search in the event consent is withdrawn. Forensic copies of computer evidence should be made as soon as possible in consent cases.

Consent Searches, continued

Third Party Consent

You can obtain consent from a third party for areas which are communal to the target and the third party. The third party must have ownership or the right to access the search area.

The search area may apply to computers when there are multiple users of a single account. However, when each user has a separate account, a third party cannot give consent or access to the target's account or file storage area. In this circumstance, the third party can only consent to a search of his or her account or any shared storage space under his or her control.

Spouses can generally consent to the search of the other spouse's property. This is true as long as the target has not asserted exclusive rights to the search area. Absent any evidence to the contrary, investigators can rely upon the consenting spouse's assertion to authority in good faith, even if the consenting spouse is later deemed to not have the authority to grant consent.

Spouses or co-tenants consent cannot be used to overcome the objections of the other tenant. The U.S. Supreme Court affirmed their position on this issue in *Georgia v Randolph* (1996) by holding that "If a potential defendant with self-interest in objecting is in fact at the door and objects in this case, the cotenant's permission does not suffice for a reasonable search, whereas the potential objector, nearby but not invited to take part in the threshold colloquy, loses out."

Parents can generally consent on behalf of juveniles living in premises under the parent's control.

Landlords generally do not have the authority to authorize a search of a rented property unless such consent is authorized under a rental agreement with the tenant.

Military commanders or magistrates can authorize the search of military facilities under their command, including the search of persons on those facilities.

System administrators may consent to the search of an entire computer or network over which they have administrative privileges.

Search Incident to Arrest

Introduction

When making an arrest or apprehending a suspect, you have the authority to conduct a search of the area under that person's immediate control for weapons and evidence of the offense. This is known as a search incident to arrest.

Relevance

An investigator should understand the authority and limitations of searches made incident to a lawful arrest or apprehension. Although you can obtain valuable evidence, there are limitations on the scope of the search.

Means, Motive, and Opportunity

There are times when the investigator may not be fully aware of the correlation between modern electronic crimes and classic Mean, Motive, and Opportunity. Here is a brief explanation of the modern definitions:

- *Means*: If the attacker has the tools (software) and the knowledge of how those tools work, they have the means.
- *Motive*: The reason that attackers commit the crime varies. It could be as simple as bragging rights, money, or political reasons.
- *Opportunity*: With so many computers on the Internet and ease of access through home networks and open wireless networks, the opportunity to commit the crime is everywhere.

Search Incident to Arrest, continued

Search Incident to Arrest

Any time you arrest or apprehend an individual, you have the authority to search his person and the area under his immediate control for weapons and evidence of the crime.

The scope of this search is an expansion of the authority under the Stop and Frisk Search in that the arresting or apprehending investigator can search for evidence of the crime. Where you can search and what you can seize are limited by the following:

- The size and possible hiding places for weapons that may threaten your safety.
- The size and nature of the evidence related to the crime for which the subject is being arrested or apprehended.
- The area under the suspect's immediate control is limited to the area in which he might reasonably reach to obtain a weapon or destroy evidence. This area has generally been limited to the room the individual is in or his surrounding area for a reasonable distance in an open area.
- Any areas open to the public or items in plain sight can be searched and seized without additional authority.

Other Search Warrant Exceptions

Immediate Threat to Life or Serious Bodily Injury

An investigator may enter and search a premise without a warrant when there is an immediate threat to life or serious bodily injury. This exception allows an investigator to come to the immediate rescue of an individual in peril. While legally inside the premises, the investigator can legally seize any evidence that may be discovered during the rescue attempt. This exception could extend to a computer or computer network if there was reason to believe that information it contained could be used to avert the loss of life or serious injury.

Immediate Threat of the Destruction of Evidence

An investigator may enter a premise without a warrant to stop the immediate destruction of evidence in a criminal investigation. The threat must be immediate. A computer or network device that contains evidence subject to being destroyed if not seized immediately could justify seizure; however a warrant should be obtained before the search occurs. Once the item is in a protected place, the exigent circumstance is vacated.

Fresh Pursuit

An investigator in pursuit of an individual may follow the individual into or through a premise. If evidence of a crime is observed during the pursuit, the investigator may legally seize the evidence. This exception is not one normally used in a digital evidence case. However the possibility exists that an officer who is lawfully in a protected place as a result of a fresh pursuit may observe contraband displayed on a computer screen. The best option in this instance would be to secure the premises and obtain a warrant. However, the officer would not be prohibited from seizing the contraband without delay.

Plain View

Any time an investigator has the right be in a physical place, any evidence of a crime visible to the investigator may be seized without a warrant and is admissible in a court of law. As discussed above, if the device seized is a computer, obtain a warrant before searching the device.

Vehicle Searches

Due to the mobile nature of motor vehicles, the U.S. Supreme Court has ruled that an investigator may search a vehicle without a warrant when the vehicle has been legally stopped and probable cause to search exists.

Other Search Warrant Exceptions, continued

Custodial Searches When a person is booked into a jail or detention facility, he/she is subject to a complete search of his/her person for weapons or contraband as well as for inventory of personal property. Evidence located during a custodial search may be legally seized without a search warrant. Storage media for digital data are becoming smaller and smaller and should not be overlooked during this search. Thumb drives and data storage cards may hold crucial evidence in a case and can be easily concealed in many locations on and in the human body.

Border Searches Individuals and vehicles arriving from foreign countries are subject to Customs inspections and search without a warrant at any port of entry into the United States. Contraband discovered during these searches may result in charges against the person attempting to bring it into the country.

This page intentionally left blank.

Lesson 2 – Internet Service Providers

Introduction

Many crimes involve the use of commercial and private networks and communications facilities. Some records associated with communications over the Internet are usually maintained by Internet Service Providers (ISPs). ISPs often maintain records of accounts, billing, transactions, and content of the communications and data that travel over their networks. At this time, ISP's are not required to keep records of traffic through their service. As an added service, many ISP's also offer data storage on their servers and e-mail services.

During an investigation, you will need to gather some or all of this pertinent information from ISPs. It is imperative that an investigator understands the proper way to request these records in order for the evidence to be admissible in a criminal proceeding.

Purpose of this Lesson

This lesson discusses the authorities under which an investigator may obtain ISP records in a criminal investigation.

Objectives

After completing this lesson, you will be able to:

- Explain which laws apply to a given authority and know where to find those laws
- Discuss the search authorities for gathering records
- Prepare requests for records

In this Lesson

The following table shows the contents of this lesson.

Topic	See Page
Legal Framework	8-20
Express Consent	8-24
Written Consent	8-26
Preservation Letters	8-29
Subpoena	8-30
Search Warrant	8-31
Available Data	8-32

Legal Framework

United States Code U.S. Code is organized in the following hierarchy:

- Titles
- Parts
- Chapters
- Sections

Parts and Chapters are generally not referenced, as the numerical sequence of the sections traverses an entire title. The section number, together with the Title number, will uniquely identify the section.

U.S. Code is normally referenced in the format:

[Title number] USC § [Section number]

The reference 18 USC § 2703 would therefore be a reference to Section 2703 in Title 18.

The Official U.S. Code is published in hard copy every six years, with annual updates being issued in between publications. The online version of the U.S. Code is accessible at:

<http://www.access.gpo.gov/uscode/index.html>

Legal Framework, continued

International Cases

With the Internet, a local agency may be faced with an investigation with an international nexus. Most of the “419 scams” or “Nigerian scams” originate in other countries. The prevalence of individuals selling illicit material online from foreign locations also poses the same challenge to local agencies.

Resolution

Local agencies can charge an individual from another country with a crime, but when the victim resides in the local agency’s jurisdiction, the agency generally cannot take custody of the individual and have them extradited to their jurisdiction for trial. As a result local agencies must look to Federal agencies that have the capacity to deal with such issues for relief.

Federal Assistance

Many federal agencies have significant experience in the area of international crime. The FBI, because of the presence of their Legal Attaches in U. S. Embassies, is most often used agency for assistance. However, many other agencies such as the U.S. Secret Service and ICE (Immigrations and Customs Enforcement) have agents in place or contacts available which can assist in the investigation.

Obviously, the amount of the loss in a financial crime must be significant to meet the threshold minimum for these agencies. Other non-economic crimes such as distribution of child pornography have different factors which determine whether the agency will pursue the investigation. If a local agency needs assistance in investigating crimes with international nexus they should contact their closest USSS Electronic Crimes Task Force or FBI Regional Computer Forensics Lab for assistance.

Another avenue of support is INTERPOL. INTERPOL is the world’s largest international member police organization that provides support and assistance in international investigations. Cyber crime is one of the topics of focus within INTERPOL as well as a number of other crimes that often have an international nexus. INTERPOL can be accessed by law enforcement agencies through their Nlets communication network. The U.S. Department of Justice is the INTERPOL National Central Bureau for the U. S.

Legal Framework, continued

Electronic Communications Privacy Act (ECPA)

Access to stored wire and electronic communications and transactional records is governed by Chapter 121 of the U.S. Code, which is currently comprised by 18 USC § 2701-2711. This chapter was enacted in 1986 by the Electronic Communications Privacy Act (ECPA).

The ECPA defines how the government can obtain stored account information from third parties. The types of information that can be obtained are broken down into three categories:

- Basic Subscriber Information - 18 U.S.C. § 2703(c)(2)
 - Name
 - Address
 - Local and long distance telephone connection records, or records of session times and durations
 - Length of service (including start date) and types of service utilized
 - Telephone or instrument number or other subscriber number or identity, including any temporarily assigned network address; and
 - Means and source of payment for such service (including any credit card or bank account number)
- Records or Other Information Pertaining to a Customer or Subscriber - 18 U.S.C. § 2703(c)(1)
 - This is a catch-all for anything else that is not content
- Contents
 - “Contents,” when used with respect to any wire, oral, or electronic communication, includes any information concerning the substance, purport, or meaning of that communication

The following table, reproduced from the U.S. DOJ publication, *Searching and Seizing Computers and Obtaining Electronic Evidence in Criminal Investigations*, summarizes the mechanisms to compel disclosure of information:

Legal Framework, continued

Item	Voluntary Disclosure Allowed		Mechanisms to Compel Disclosure	
	Public Provider	Non-Public Provider	Public Provider	Non-Public Provider
Basic subscriber, session and billing information	Not to government, unless § 2702(c) exception applies [§ 2702(a)(3)]	Yes [§ 2702(a)(3)]	Subpoena; 2703(d) order; or search warrant [§ 2703(c)(2)]	Subpoena; 2703(d) order; or search warrant [§ 2703(c)(2)]
Other transactional and account records	Not to government, unless § 2702(c) exception applies [§ 2702(a)(3)]	Yes [§ 2702(a)(3)]	2703(d) order or search warrant [§ 2703(c)(1)]	2703(d) order or search warrant [§ 2703(c)(1)]
Accessed communications (opened e-mail and voice mail) left with provider and stored files	No, unless § 2702(b) exception applies [§ 2702(a)(2)]	Yes [§ 2702(a)(2)]	Subpoena with notice; 2703(d) order with notice; or search warrant [§ 2703(b)]	Subpoena; ECPA doesn't apply [§ 2711(2)]
Unretrieved communication, including e-mail and voice mail (in electronic storage more than 180 days)	No, unless § 2702(b) exception applies [§ 2702(a)(1)]	Yes [§ 2702(a)(1)]	Subpoena with notice; 2703(d) order with notice; or search warrant [§ 2703(a,b)]	Subpoena with notice; 2703(d) order with notice; or search warrant [§ 2703(a,b)]
Unretrieved communication, including e-mail and voice mail (in electronic storage 180 days or less)	No, unless § 2702(b) exception applies [§ 2702(a)(1)]	Yes [§ 2702(a)(1)]	Search warrant [§ 2703(a)]	Search warrant [§ 2703(a)]

Express Consent

Express Consent Doctrine

Express Consent is consent that is derived from an individual's actions as the result of documents and notices provided to that individual before an incident occurs. This type of consent is usually seen in logon banners or signs advising that use of the system or entry onto government property "implies" consent to be monitored or searched. The act of entry onto the system or property constitutes informed, voluntary consent.

Relevance

Whenever a cyber crime is committed that involves the unauthorized use of a computer or network, the investigator should establish whether or not authority to gather information exists under the Express Consent. By doing this, the investigator can establish the initial basis of authority for the monitoring of an individual's actions or the search of the individual and items under his or her control.

Establishing Authority and Ownership

During the initial investigation of a cyber crime, the investigator should firmly establish the owner or Designated Authorization Authority of any computer or network involved in the investigation. This ownership or authority should be documented in the investigation report and verified by supporting documentation in the form of policies, orders, copies of ownership records or written statements.

Express Consent, continued

Establishing and Documenting Express Consent

When a user logs onto a computer or network, there should be an initial warning banner that explains the authority to access the computer or network as well as any implications such access may have for the user. Here is an example of a common warning banner:

THIS IS A STATE OF MARYLAND COMPUTER SYSTEM

This computer system including all related equipment networks and network devices (specifically including Internet access) are provided only for authorized State of Maryland use. State of Maryland computer systems may be monitored for all lawful purposes to ensure that their use is authorized for management of the system, to facilitate protection against unauthorized access and to verify security procedures survivability and operational security. Monitoring includes active attacks by authorized State of Maryland entities to test or verify the security of the system. During monitoring, information may be examined, recorded, copied and used for authorized purposes. All information, including personal information, placed on or sent over this system may be monitored. Use of this State of Maryland computer system, authorized or unauthorized, constitutes consent to monitoring of this system. Unauthorized use may subject you to criminal prosecution. Evidence of unauthorized use collected during monitoring may be used for administrative, criminal or adverse action.

The banner should require some type of action by the user to acknowledge the presence of the banner. Banners that flash and disappear without user interaction may not suffice to establish express consent.

For example, anyone who logs onto a State of Maryland computer system or network must click through a banner screen before proceeding to the logon prompt. Once users click through, they have consented to monitoring of their activity, and access to their data on the network or individual computer.

Written Consent

Introduction

Fourth Amendment rights, like other constitutional rights, may be waived, and a person may consent to a search of his person or premises. The U.S. Supreme Court, however, has insisted that the burden is on the prosecution to prove that the consent was voluntary and the person was aware of the right of choice.

Relevance

Investigators should understand how to obtain voluntary consent and the limits of that consent. Although verbal consent to search is valid, it is very important that the consent be documented in writing to avoid issues later at trial. Written consent from the account holder may be necessary to obtain records from banks, hospitals, or similar businesses unless the investigator has some other legal authority to obtain them like a subpoena or search warrant.

Written Consent, continued

Written Consent

Whenever an investigator establishes that an individual has ownership or other legal authority over any item or area, the investigator can request consent to search the individual and property under his or her control. Whenever possible, you should obtain this consent in writing and establish the following in the consent document:

- The person who gives the consent has ownership or authority over the premises or, in the case of a cyber crime, the computer system to be searched.
- The consent must be freely given.
- The consent must not have been coerced. Actual knowledge of the right to refuse consent is not essential to the issue of voluntary consent, and therefore police are not required to acquaint a person with his rights, as through a Fourth Amendment version of Article 31. Giving the individual a warning about his or her rights concerning consent has been taken by the courts as further indication of voluntary consent. Consent will not be regarded as voluntary when the officer asserts his official status and the individual yields to that authority rather than make his or her own determination.
- Consent may be withdrawn by the person giving consent at any time.
- The scope of the consent search can be limited by the person giving consent
- The investigator should document all of the conditions under which the consent is being given, particularly those listed above and obtain the signature of the person giving consent on the document.

Written Consent, continued

Undercover Deception

When consent is obtained through the deception of an undercover officer or an informer gaining admission without advising an individual who he is, the U.S. Supreme Court has held that the individual has simply assumed the risk that an invitee would betray him, and evidence obtained through the deception is admissible.

This exception rarely applies to the production of records unless the individual is a business entity in possession of the records and the records are obtained within the course of an undercover investigation.

In these cases, obtaining written consent may not be possible. The investigator must carefully document his or her search authority through written reports or other means such as undercover recordings.

Third Party Consent

Additional issues arise in determining the validity of consent to search when consent is given not by the individual suspected of the offense, but by a third party. In the earlier cases, third party consent was considered sufficient if that party possessed common authority over or other sufficient relationship to the premises or effects sought to be inspected.

For example, spouses are presumed to have authority to consent to the search of the other spouse's property unless there is evidence to the contrary. The best policy is to establish spousal control and access in writing.

Juveniles

The law concerning the authority of a juvenile or minor to intelligently give consent varies from jurisdiction to jurisdiction. It is generally accepted that the parent retains the consent authority for a juvenile and that consent should not be obtained from minors.

Preservation Letters

Introduction

Storage and destruction of electronic logs and records is much more dynamic than traditional methods of recordkeeping. It takes time to obtain court orders (up to and including warrants), but there are no federal laws that require companies to preserve electronic records for any amount of time. Some companies maintain records and logs for many months; some maintain no records at all. This can present a problem for an investigator.

Relevance

What if pertinent records for your investigation exist at an Internet Service Provider company today, but you cannot get a subpoena for three days? Will the records you need still be there?

The answer is that it depends on the age of the data and the company's retention practices. The data may not be there when you need it.

Part of the ECPA is designed to ensure that available data is not lost during the time it takes to obtain the necessary order.

Issuing a Preservation Request

18 USC § 2703(f)

The benefit of the preservation request is that an investigator can issue it quickly and directly to preserve information for 90 days. Preservation requests are designed to ensure that the specified information will still be there when the appropriate legal process is served to obtain it.

18 USC § 2703(f)(1) states:

A provider of wire or electronic communication service or a remote computing service, upon the request of a governmental entity, shall take all necessary steps to preserve records and other evidence in its possession pending the issuance of a court order or other process.

There is no required format for § 2703(f) requests, but most agencies have developed their own preferred format over time, and it is usually in the form of a letter. Technically, the request can be verbal.

Subpoena

Introduction

The subpoena is a court order requiring a person or business entity to produce records or testimony. Failure to comply with a subpoena may result in penalties or criminal charges.

Relevance

The subpoena is commonly used by investigators and prosecutors to require an individual or business to produce records or testimony. Subpoenas are used to obtain stored transactional records (basic subscriber information) and in some circumstances, may even be used to obtain stored wire and electronic communications (content).

Obtaining a Subpoena

In a cyber crime investigation, the investigator can request that the prosecutor issue a subpoena for records that are under the control of an individual or business. The prosecutor has the authority to issue the subpoena and have the person or business served with the document. Once served, the individual or business has the right to argue before a competent court of jurisdiction why the records should not be produced.

Subpoenas are generally used to obtain records from individuals or businesses that are not the target of the investigation. In cases where the individual or business is the target of the investigation, the search warrant is preferred to prevent the destruction or alteration of the records being sought.

Search Warrant

Introduction

While most information stored by an ISP can be retrieved with a subpoena or a “D” order which is discussed next, at least one type of information requires a search warrant under [§ 2703(a)].

Relevance

A search warrant is required to seize unretrieved communication, including e-mail and voice mail (in electronic storage 180 days or less) from an Internet Service Provider. Until 2006 opened mail or mail that had been accessed by the user, but still stored at the ISP over 180 days could be obtained without a search warrant. In *Warshak v U.S.* the 6th Circuit Court of Appeals ruled that the provisions of 18 U.S.C. § 2703, part of the Stored Communications Act, which allowed law enforcement to seize stored communications over 180 days old without a warrant and without notifying the subscriber and allowing them to bring their objections before the court violated the 4th amendment of the U.S. Constitution.

Obtaining a Search Warrant

Search warrants issued under 2703(a) for information held by an Internet Service Provider require the same standard of probable cause discussed earlier. Generally this is accomplished by the applicant preparing an affidavit which contains the facts known to him which lead him to believe that the ISP holds relevant evidence in the case. The affiant must swear to the contents of the affidavit before the court will, if persuaded by the affidavit, sign the warrant.

Available Data

Introduction

No current U. S. law requires that Internet Service Providers maintain information of traffic through their service. All of the U. S.-based major ISP's do maintain logs of information on traffic as well as connection logs, however the length of time this information is are kept and what is kept varies.

Timeliness

When you need information from an ISP that is critical to your case, it is very important that you send them a preservation letter as outlined in a previous lesson. Failure to do so may result in the information not being there when you get the appropriate legal authority to obtain the information.

Available Data

Be aware that some small ISP's do not store any logs pertaining to traffic through their offices. They generally will have some sort of subscriber information because someone is paying the bill and they need a way of tracking the payments. Even if the subscriber information is bogus, the credit card number that is paying the bill may lead you back to the real subscriber.

Data that you may be able to obtain from an ISP include:

- Subscriber information to include payment information
- IP's addresses used and date and times used
- Session lengths (Connection Logs)
- Email- opened and unopened
- Stored Files (a service offered by some ISP's)
- Related account information (has multiple accounts)
- Phone numbers used to access using dial up connections

Module 9

Fundamentals of Log Analysis

Overview The analysis of computer network intrusions is a difficult task. The Scientific Method provides a general framework that can be used to effectively guide the investigation.

Purpose of this Module This module provides a review of computer intrusion methods and a description of the Scientific Method as it applies to intrusion investigation. This description is focused on the discovery and analysis of log-based artifacts.

Objectives After completing this module, you will be able to:

- Describe the main steps of the Scientific Method
- Explain how the Scientific Method can be applied to digital forensic analysis
- Use the initial observations in a case to determine the most likely location of additional, related artifacts
- Apply the analysis techniques learned in the previous modules to analyze log files that contain evidence of an intrusion

In this Module The following table shows the contents of this module.

Topic	See Page
Lesson 1 – Understanding Network Traffic	9-3
Lesson 2 – The Scientific Method and Intrusion Analysis	9-9
Lesson 3 – Observing Intrusion-related Activity and Generating a Hypothesis	9-15
Lesson 4 – Predicting the Nature and Location of Intrusion Artifacts	9-25
Lesson 5 – Using Log Analysis to Evaluate an Intrusion Hypothesis	9-37

This page intentionally left blank.

Lesson 1 – Understanding Network Traffic

Introduction

The first step in learning how to analyze log files is to look at the types of data traffic you will typically see on a network. How these protocols pass from system to system is important to the investigative and analytical process.

Purpose of this Lesson

The purpose of this lesson is to describe the network traffic and to discuss how it can be used to guide an intrusion investigation.

Objectives

After completing this lesson, you will be able to:

- Define the different types of network traffic
- Recognize which types of traffic are of interest in intrusion investigations.

In this Lesson

The following table shows the contents of this lesson.

Topic	See Page
Overview of Network Traffic	9-4
Investigation Techniques	9-5

Overview of Network Traffic

The Internet

When you look at the Internet or any network that connects two or more computers you will see common communications types. TCP/IP is the language of the network and there are a number of services which use TCP/IP to communicate.

For example, we will look at a basic example of the service HyperText Transfer Protocol (HTTP) and see how it actually works.

Getting from One Place to Another

For HTTP or any other service that generates network traffic to work there must be one system that is serving or hosting information and another system that is requesting the service data. In the case of HTTP, there is both a Web server and a Web client connected to a network. The network traffic that exists when the client or browser requests a Web page from the server proceeds in this manner:

- 1) *SYN*: The client sends a Synchronize packet to the server, beginning the three-way handshake which starts the conversation.
- 2) *SYN-ACK*: The server sends a Synchronization Acknowledgement, acknowledging the start of the conversation.
- 3) *ACK*: The client then sends an acknowledgement to the server completing the three-way handshake. The conversation is started at this point.
- 4) *GET*: The client requests a page from the server. If this is a general request to a web site like `www.somewhere.com` then the GET request is for the web root document indicated with a backslash / after the get command. Otherwise the name of the page will be part of the get request, such as `faq.html`.
- 5) *200 OK*: The server will send a response that includes the status code for the page requested. Usually the code 200 OK is sent, meaning that the page was found and will immediately follow:
 - If the specific page requested is not recognized by the server the now famous 404 Page not found is displayed in the browser.

The page will now be transferred to the browser program and the page will be displayed on the client system.

Overview of Network Traffic, continued

Ports in a Data Storm

There is an additional bit of information that follows every IP address used on the network: the port number. There are 65535 port numbers available on most computer systems. The Internet Assigned Numbers Authority (IANA) has the role of assigning the types of traffic to each of these port numbers. This is done so that programmers can agree on which ports are used for specific types of network traffic. HTTP traffic is assigned port 80.

Types of Ports

There are three types of port numbers, Well Known Ports, Registered Ports and Dynamic Ports.

Well Known Ports are the numbers ranging from 0 to 1023.
Registered Ports are the numbers ranging from 1024 to 49151
Dynamic Ports are the numbers ranging from 49152 to 65535

Documentation on the current assignment of these numbers can be viewed by going to www.iana.org/assignment/port-numbers

Common Ports and Assigned Traffic

Below are some of the most commonly used ports and the traffic types assigned to them in the Well Known Port range:

- (20) File Transfer Protocol (FTP)
- (21) File Transfer Control (FTP)
- (22) Secure Shell Remote Login
- (23) Telnet
- (25) Simple Mail Transfer Protocol (SMTP E-mail)
- (53) Domain Name Service (DNS)
- (80) HTTP (Web)
- (443) Secure Socket Layer (HTTPS)

Below are some of the common assigned ports in the Registered Port Number range:

- (1025) Network Blackjack
- (1080) SOCKS
- (1169) Tripwire
- (1214) KAZAA
- (1433) Microsoft SQL Server
- (1689) Firefox

The Dynamic Ports can be used by any service or protocol at any time depending on the random assignment used by the operating system for a given computer.

Investigation Techniques

Introduction

When investigating network traffic, one of the first things to look for is traffic types that are on the wrong assigned ports. For example, when most government agencies realized that employees were using AOL Instant Messenger at work, the network administrators closed the AOL IM port 531. Users quickly discovered this and changed their clients to use the unblocked HTTP port 80.

Looking at a packet capture file and seeing instant message traffic on port 80 is an indication of the sophistication of the end user. They are technologically aware enough to know that changing the port number used by a service or program will circumvent the network security profile.

Types of Traffic to Watch For

Here are some of the types of traffic you will probably see during your network investigations and the potential issues you might look for.

HTTP (port 80) – Because most firewalls and routers will pass traffic on port 80, it is a popular port for malicious code transfer or for communication of other protocols that have been blocked. In some cases, you will see programs that have opened backdoors on systems transferring information on this port. Advanced attackers will embed malicious information in HTTP packets hoping that firewalls and intrusion detection systems will pass the information.

E-mail (port 25) – Although not as common a port for non-e-mail traffic, this port is one worth watching simply because so many attacks originate in e-mail messages.

USENET/NNTP (port 119) – This is an important protocol for law enforcement to watch since many newsgroups are used for distribution of pornography in all forms. This protocol is still used as a way to transfer bootleg software, movies, music and other copyrighted material.

Investigation Techniques, continued

Types of Traffic to Watch For, continued

Internet Relay Chat (ports 6666-6669) – IRC is another protocol that is used heavily for Peer-to-Peer transfer of copyrighted and illicit materials. Malware and Botnet traffic is seen on these ports as well.

File Transfer Protocol (port 21) – FTP is used for transferring files therefore if your case may involve transfers of illicit information of any kind it would make sense to monitor FTP traffic. There are a number of malware attacks against FTP ports as well.

Peer-to-Peer (Any ports) – P2P protocols are some of the hardest to monitor and investigate because of the nature of the protocol. The two systems that are transferring information can use any port they agree on and the transfer of a file may actually take place between multiple systems at once. This can make the reassembly of transferred files extremely difficult. Tools like Wireshark will typically indicate that P2P traffic is taking place.

Baselines

One popular way to make network traffic analysis easier is the baseline method. By taking packet capture snapshots of normal network traffic and then comparing suspicious traffic captures to that baseline you can more quickly determine where the investigation should focus.

This page intentionally left blank.

Lesson 2 – The Scientific Method and Intrusion Analysis

Introduction

The Scientific Method is used as a guide for investigating any problem, including a network intrusion. It is a simple but effective process by which you generate a hypothesis based upon observed events, then design and select analysis tasks to help you evaluate that hypothesis.

Purpose of this Lesson

The purpose of this lesson is to describe the Scientific Method and to discuss how it can be used to guide an intrusion investigation.

Objectives

After completing this lesson, you will be able to:

- Define the Scientific Method
- Explain how the Scientific Method can guide an intrusion investigation

In this Lesson

The following table shows the contents of this lesson.

Topic	See Page
Overview of the Scientific Method	9-10
Digital Forensic Analysis and the Scientific Method	9-12

Overview of the Scientific Method

Definition

The *Scientific Method* is a process for investigating a set of observations. The method is conducted by formulating a hypothesis about observed events that are of interest, then using deductive/inductive logic to formulate processes for evaluating that hypothesis. The developed processes are then carried out and their results are used to support, contradict, or modify the hypothesis.

Steps

The following steps comprise the Scientific Method as it is used in this course:

1. *Observation*: Observing one or more events or sets of events. Observation establishes the facts surrounding these events to identify their cause and consequences.
2. *Hypothesis*: A hypothesis is generated that explains the observed events, including their root cause, interrelationship, and consequences.
3. *Prediction*: Predictions are made as to the possible nature and location of artifacts in the evidence that will either support or contradict the hypothesis.
4. *Evaluation*: Performing procedures that test for the presence of artifacts that support, falsify, or modify the hypothesis.
5. *Conclusion*: Formation of a conclusion, based upon the results of tests performed during the Evaluation step. The conclusion states one of the following:
 - The hypothesis is supported by the facts
 - The hypothesis is contradicted by the facts
 - The facts indicate that a new or modified hypothesis should be constructed due to new observations or lack of relevant results.

Overview of the Scientific Method, continued

Additional Characteristics

The Scientific Method has principles that are not specific steps in the method, but factors vital to ensuring the Scientific Method is carried out properly.

- *Repeatable*: All evaluations and tests conducted during an iteration of the method should be repeatable. This is to ensure that results can be verified by others who want to test them for mistakes, confounding variables, spurious relationships, etc.
- *Cyclic*: The Scientific Method is cyclic, meaning that a scientist may need to perform many iterations of the method, test and resting a hypothesis, or generate additional hypothesis to finally gain a clear understanding of the originally observed events.
- *Empirical*: All evidence used in the hypothesis must be based on or derived from observation rather than pure reasoning, faith, common sense, etc.
- *Falsifiable*: A hypothesis that is established and tested using the Scientific Method should be falsifiable. In other words, there should be a way to test for contradicting evidence as well as supporting evidence.
- *Objectivity*: Observations and the results of any evaluations must be interpreted as objectively as possible.

Variances in the Use of the Scientific Method

Not all fields of inquiry use the same steps for the Scientific Method and the names of the steps can differ. So when researching the method you will sometimes encounter different formats within different reference sources. An implementation of the Scientific Method is valid so long as it follows the principles outlined above.

Digital Forensic Analysis and the Scientific Method

The Use of the Scientific Method During Digital Forensic Analysis

The Scientific Method provides a useful guide when attempting to locate items of interest within digital media, or copies of digital media. It is also useful for the incident responder when he/she attempts to identify devices that may contain information related to a series of events.

Example

As an example, an incident responder investigates the appearance of several IDS alerts indicating an attack against a Web server. These alerts would be the initial *observations*. The incident responder might then form a *hypothesis* that the Web server had been attacked and compromised by the method indicated in the alerts.

To test this hypothesis, the analyst would then deduce (*predict*) the most probable location of artifacts that would support or contradict the hypothesis that the system had been successfully attacked. Supporting artifacts might include unauthorized Registry entries, the presence of malicious code, additional IDS alerts, unauthorized user accounts, and so forth.

Contradicting artifacts could be other log entries that show that the observed events are part of normal activity for an application. The analyst would gather data from devices that contain these artifacts, and *evaluate* that data for their presence.

The examiner finds artifacts that support the hypothesis that the system was successfully attacked. The examiner may then *conclude* that the hypothesis was correct, and proceed to write a report.

As an alternative, the hypothesis may have been falsified due to the discovery of artifacts indicating a legitimate technical reason for the IDS alerts occurrence such as a standard false positive. In addition, the investigator may not find sufficient evidence to make any conclusion about the hypothesis, in which case he/she may create a new hypothesis.

Digital Forensic Analysis and the Scientific Method

The Use of the Scientific Method for Computer Intrusion Investigations

Computer network intrusions can be complex and difficult to track down. In an enterprise environment, an attack can span multiple networks that include thousands of computer systems. One danger for the investigative team is that they will spend too much time acquiring and analyzing data from unrelated systems.

The Scientific Method helps avoid this pitfall by encouraging you to follow a logical process to determine how to conduct an investigation. The key element is the link between observed events and subsequent investigative tasks.

By creating hypotheses based on real events, you are more likely to perform analysis tasks that produce results, and less likely to follow unproductive tangents. Subsequent lessons in this module show you how to apply the method to intrusion analysis.

This page intentionally left blank.

Lesson 3 – Observing Intrusion-related Activity and Forming a Hypothesis

Introduction

The first step of the Scientific Method applied to an intrusion is to identify the current set of observations and form a hypothesis based upon those observations.

Purpose of this Lesson

The purpose of this lesson is to learn the common types of intrusion-related observations and how to form a hypothesis based upon them.

Objectives

After completing this lesson, you will be able to:

- Describe common intrusion-related observations
- Form a hypothesis
- Describe common incident classifications

In this Lesson

The following table shows the contents of this lesson.

Topic	See Page
Common Observations	9-16
Hypothesis Formation	9-19
Incident Classification	9-21

Common Observations

Observations and Network Intrusions

Network intrusion investigations should normally begin with one or more specific observations. These observations guide the formation of a hypothesis as to what may have occurred.

Common Primary Observations

Many different events can spark an intrusion investigation. Some examples include:

- *Antivirus alerts:* AV systems will sometimes notice the presence of one or more malicious files. This is a common intrusion indicator, especially when trojans, backdoors, and rootkits are detected.
- *IDS/IPS alerts:* Intrusion detection system alerts are messages specific to attack-related activity and are common first warning events.
- *System/applications errors:* Compromised systems will sometimes experience errors due to problems caused by attack-related activity. Attacks against applications can cause those applications to crash. If administrators are unable to find legitimate reasons for crashes during their initial troubleshooting, then those crashes may be indicators of an intrusion.
- *Abnormal authentication patterns:* Repeated failed authentication attempts, unusual login times or attempts to authenticate as a non-existent user account are indicators of an attempted attack.
- *Access control list violations:* Failed attempts to communicate through a barrier system such as a firewall or proxy server that is logged can be an indicator that an attempt is being made to breach a network.
- *Generic unusual activity:* Sometimes the initial observation is simply something that a user or system administrator noticed as being abnormal and reported to the designated security contact. A common abnormality is activity occurring at an unusual time that would otherwise appear legitimate, such as file transfers.

Common Observations, continued

Supplementary Observations

The incident responder should make supplementary observations before creating a hypothesis. These are not directly observed events, but rather sets of data that the responder should collect in any security incident. Examples of this type of data includes:

- Network diagrams: Logical and physical diagrams of the networks where the event occurred.
- Device documentation: Lists of device names and configuration data. This information is especially vital for devices directly involved in observed events.
- Contact information: Names, phone numbers, e-mail addresses, etc. for witnesses and people responsible for the affected networks and systems.
- Other data: Any other details regarding the affected devices and networks that may seem pertinent.

Common Observations, continued

Common Observation Attributes

Observations made during network intrusions will have attributes that should be recorded. These attributes should be gathered correctly from the incident responder or the network administrator on site. These attributes include, but are not limited to the following:

- **Date/time:** Record when the event occurred, as well as its duration.
- **IP addresses:** If the event is a log entry that includes an IP address, or if it involves a system with an IP address, then that IP address should be recorded.
- **Port numbers:** If the event is a log entry that includes port numbers, or it involves an application that engages in network communication over a specific port, then that port should be recorded.
- **Accounts and aliases:** If the event involves a specific user account or alias, then that name should be recorded, as well as the name of the specific individual that uses that account or alias, if that information is known.
- **Host names and aliases:** The host names and aliases for any system involved in an event should be recorded.
- **Files:** At a minimum, the name and full path for any files involved in an event should be recorded. If available, other useful attributes that can be recorded about a file include hash value and file system date/time stamps (created, modified, etc.)
- **General description:** A general description as to the nature of each event should also be recorded.

Recording Observations

Observations can be recorded in many different forms including written notes, office documents, and databases. You should use the approved and tested method used by your organization. This course uses a spreadsheet template for recording this data.

Hypothesis Formation

Hypothesis Formation

The initial set of observations from an incident will enable you to form a hypothesis regarding the incident. This hypothesis should include a statement regarding each of the following:

- “What/How:” Basic description of the main event(s). This may include a common incident classification
- “Where:” List the known and probable physical locations and network segment locations of the incident.
- “When:” List the known and probable timeframe of the incident.
- “Who:” List identifying information for the individual(s)/computer(s) known to be involved or likely involved in the incident.
- “Why:” List the most likely motive(s).

At the start of an investigation, many details are still unknown. Therefore the initial hypothesis may be broad. As the investigation proceeds and more facts are discovered, multiple cycles through the Scientific Method may yield more specific hypothesis.

Example Hypothesis

A hypothesis statement can be recorded easily in table form as noted below.

Category	Statement
What	A <incident classification> occurred against <Victim System(s)>, resulting in <Resultant access, theft or damage>.
Where	ABC Corp., Reston VA <Address>
When	First Related Event: 8/16/07 0715 EST Last Related Event: 8/19/07 1611 EST
Who	Attacker(s) Name/Alias: <Name or Handle> Attacker System(s) Hostname: <Hostname> Attacker System(s) IP: <IP> Victim System Hostname(s): <Hostname> Victim System IP(s): <IP>
Why	Possible reasons for the <incident classification> to have taken place.

Hypothesis Formation, continued

Multiple Hypotheses

At some point in the investigation, you may decide that the incident is too large and complex for a single hypothesis. You may then need to establish multiple hypotheses to account for different parts of the incident.

For instance, a large enterprise intrusion may have signs that the attacker entered through a public Web server and through several compromised workstations. To effectively pursue each possibility, you might create one hypothesis to pursue each potential method of entry.

If you are a manager or lead investigator, you may assign different investigators to separately investigate each hypothesized method of entry. You could even create a third hypothesis to account for how the attacker(s) are extracting stolen data from the network.

There is no rule for determining how many hypothesis to create or how detailed they should be. Hypotheses should reflect the size and complexity of the incident.

Incident Classification

Incident Classifications

You should implement an incident classification schema to ensure a common vocabulary between you and the organization requesting assistance. The classification should be broad enough to capture the major types of incidents you might encounter. Here are the recommended incident classifications:

- Denial of Service
- Malicious Code
- Unauthorized Access
- Inappropriate usage
- Suspicious activity
- Multiple Component
- Other

A further description of each of these is provided below, including lists of common observations that may lead to you to include the classification in your hypothesis.

Denial of Service

Denial of service is an attack that prevents or impairs the authorized use of networks, systems, or applications. Observations that could lead to this classification include:

- A network service is unavailable for an unknown reason
- A computer network is saturated with an excessive amount of network traffic
- An application is saturated with authentication or service requests
- A application or operating system is not functioning for an unknown reason

Malicious Code

Malicious code is any computer program or group of programs that perform undesirable activity on a system. Observations that could lead to this classification include:

- Antivirus alerts
- IDS alerts that indicate malicious code
- A higher than normal volume of network traffic
- Computer systems crash or malfunction for an unknown reason
- Egress communication not initiated by a user or an authorized application

Incident Classification, continued

Unauthorized Access

With *unauthorized access*, a person gains logical or physical access without permission to a network, system, application, data, or other resource. Observations that may lead to this classification include:

- User account authentication at abnormal times, or at times where the user to which the account was assigned denies having been on the subject system
- Presence of unauthorized user accounts
- Missing data
- Logged data access at abnormal times or by a user account not normally used for such access
- Presence of unauthorized computer programs
- Presence of large archives (TAR, RAR, Zip, etc.) of data files for which there is no explanation
- Common observations from any other type of intrusion-related activity

Inappropriate Usage

With *inappropriate usage*, a person violates acceptable computing use policies. Observations that may lead to this classification include:

- Web browsing sessions to websites containing unauthorized workplace viewing material
- Inappropriate e-mails sent to coworkers or from a work account
- Recorded network traffic that indicates the presence of an unauthorized application, such as a peer-to-peer file sharing application

Suspicious Activity

With suspicious activity, the security operations personnel notice unusual activity not specifically related to a known threat, but in their experience with the current environment is unexplainable. Observations that may lead to this classification include:

- Increase network activity
- Increase CPU activity on a system
- Unexplained network activity

Incident Classifications, continued

Multiple Component

The multiple component classification has a single incident that encompasses two or more incidents. For example, a malicious code infection leads to unauthorized access to a host, which is then used to gain unauthorized access to additional hosts. Examples include the following:

- Workstation affected by a virus and scanning the network
- Server relaying IRC traffic

Other

The category “Other” serves as a catch all group for newly identified Exploits that do not fit in any of the previously listed categories. Examples include the following:

- Penetration Testing
- Innovative ways to attack a system
- Zero-day Exploits

This page intentionally left blank.

Lesson 4 – Predicting the Nature and Location of Intrusion Artifacts

Introduction

Once you develop a hypothesis, you can use it along with the observed events to determine the most likely location(s) of any supporting or contradicting artifacts.

Purpose of this Lesson

The purpose of this lesson is to teach you how to determine potential locations of artifacts related to your hypothesis.

Objectives

After completing this lesson, you will be able to:

- Determine the applications and network traffic types that were involved in observed events
- Determine the flow of network traffic related to observed events
- Predict artifact location based upon the network architecture, probably traffic flow and related applications

In this Lesson

The following table shows the contents of this lesson.

Topic	See Page
Predicting the Nature and Location of Intrusion Artifacts	9-26
Relating Observed Events to Network Services and Traffic Types	9-27
Mapping Observed Activity to Traffic Flow	9-29
Using Traffic Flow and Service Type to Predict Artifact Location	9-33

Predicting the Nature and Location of Intrusion Artifacts

Finding Intrusion Artifacts

Finding artifacts related to a network intrusion can be a difficult process due to the vast amount of data in which these artifacts may reside. To stay focused, use the current hypotheses for the investigation to guide your search. This is done by:

- 1) Mapping observed events to related applications and traffic types.
- 2) Map observed activity to traffic flow (preferably using an accurate network diagram) so that you know which network path related to network traffic may have taken.
- 3) Using the probable traffic flow, involved applications and traffic types, determine which specific devices may have artifacts of the observed and hypothesized events.
- 4) Establish a plan for gathering data from the identified devices, and for identifying any relevant artifacts within those data sets.

The following sections of this lesson will cover these tasks in more detail.

Relating Observed Events to Applications and Network Traffic Types

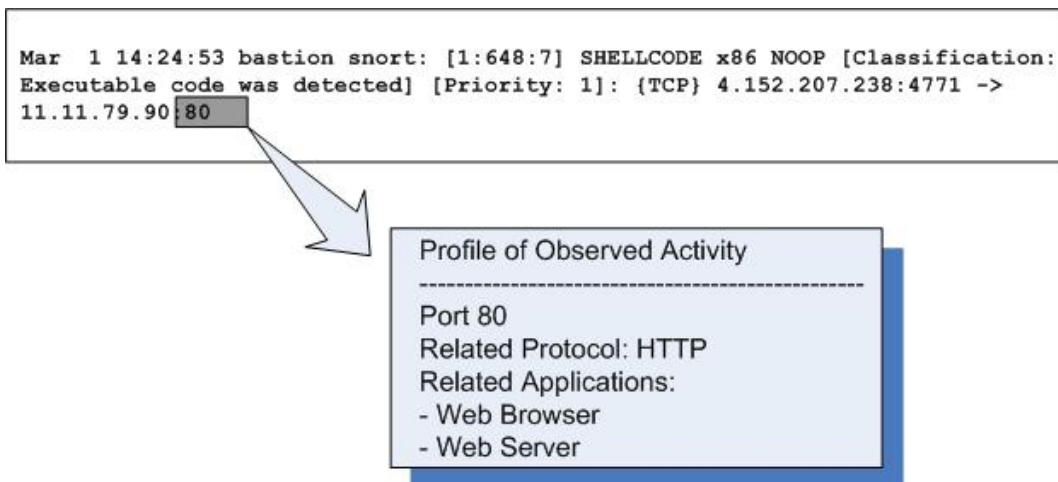
Relating Observed Events to Applications

You need to correlate all observed events to the applications involved. This will help you to locate potential artifacts. For instance, if the observed event was a buffer overflow IDS alert with a destination port of 80, you could surmise that the target application of that attack may be a Web server such as IIS or Apache. Recognizing this, you would place this application on your list of potential artifact sources and gather and analyze the logs from that application.

At a minimum, you should perform the following tasks to help identify involved applications:

- Identify network traffic types that correspond to observed TCP/UDP ports. For instance, observed TCP port 25 traffic indicates that SMTP is most likely involved.
- Identify applications related to observed and/or extrapolated network traffic types. From the example above, if SMTP were the likely protocol, then that would indicate that an e-mail server and client application were probably also involved.

The image below illustrates the concept of profiling an event.



Relating Observed Events to Applications and Network Traffic Types, continued

Additional Applications

Also list applications that meet the following criteria:

- Identify applications that have the capability of logging activity related to network traffic types and applications you have already singled out. For instance, SMTP gateways would have the capability of logging data about traffic between e-mail servers and clients.
- Identify applications directly involved in the generation of observed events. This also includes security devices/applications that produced log files that contained initial observations. Following the example from the previous page, the Snort IDS that generated the alert would be added to the list of applications that may contain relevant artifacts.

```
Mar  1 14:24:53 bastion snort: [1:648:7] SHELLCODE x86 NOOP [Classification: Executable code was detected] [Priority: 1]: {TCP} 4.152.207.238:4771 -> 11.11.79.90:80
```

Profile of Observed Activity

Port 80
Related Protocol: HTTP
Related Applications:
- Web Browser
- Web Server
- **Snort IDS**

Recording Applications and Network Traffic Types

The data produced here can be kept in any form with which you are comfortable. You could keep a list of potentially involved applications and network protocols in your notes, a database, or spreadsheet.

Mapping Observed Activity to Network Traffic Flow

Network Traffic Flow and Intrusion Artifacts

One simple way to identify devices that may contain relevant data is to locate all devices that related traffic may have passed through. For instance, if the investigator believes that intrusion-related traffic passed through a specific point of ingress/egress for the network, you can surmise that the devices at that point (firewalls, routers, IDS sensors, etc.) could potentially contain important artifacts.

Mapping Observed Activity to Network Traffic Flow

Enterprise networks can be very large and complex. Observe events to determine probable routes for related traffic. You will need the following items to map traffic flow:

- A logical or physical network diagram, and/or access to a network administrator that has working knowledge of the current topology. This diagram should be broad enough to include all points of ingress/egress from the affected network segments, including paths to the Internet.
- IP addresses for devices potentially involved in the incident
- Ports and protocols corresponding to related network protocols

With these items, identify all routes between the affected devices and between those devices and the Internet. Record these routes in your notes, or mark them on working copies of any network diagrams you were able to obtain.

Mapping Observed Activity to Network Traffic Flow, continued

Other Routes of Interest

In addition to the main routes of network traffic between affected devices and the Internet, the following routes may also be of interest:

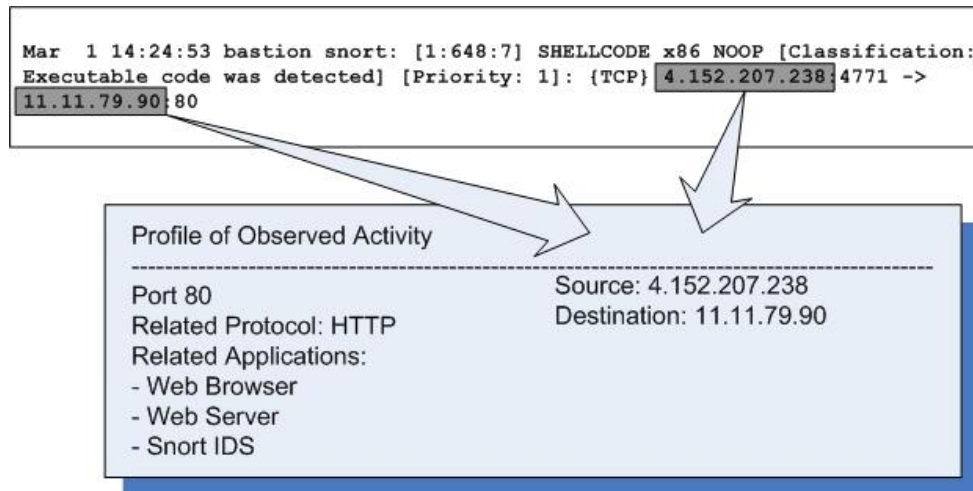
- Alternate points of network traffic ingress/egress from the network segment on which each device is resident.
- Identify any routes to major service network segments that are not inline with the default gateway (network segments with directory servers, e-mail servers, file and print servers, backup servers etc.).
- Routes used by incoming traffic to the affected network segment, if not the same as the default outbound route (routes used by public service requests to the segment, internal service requests, VPN pathways, etc.).
- In the network, check specific protocols sent through alternate routes to reach proxy servers. If some are found, identify the routes between the affected network segment and those proxy servers.

The devices along the routes identified in the questions above are all potentially in scope. Use the additional criteria listed on the following page for determining how to prioritize devices for acquisition.

Mapping Observed Activity to Network Traffic Flow, continued

Example: Adding Source and Destination

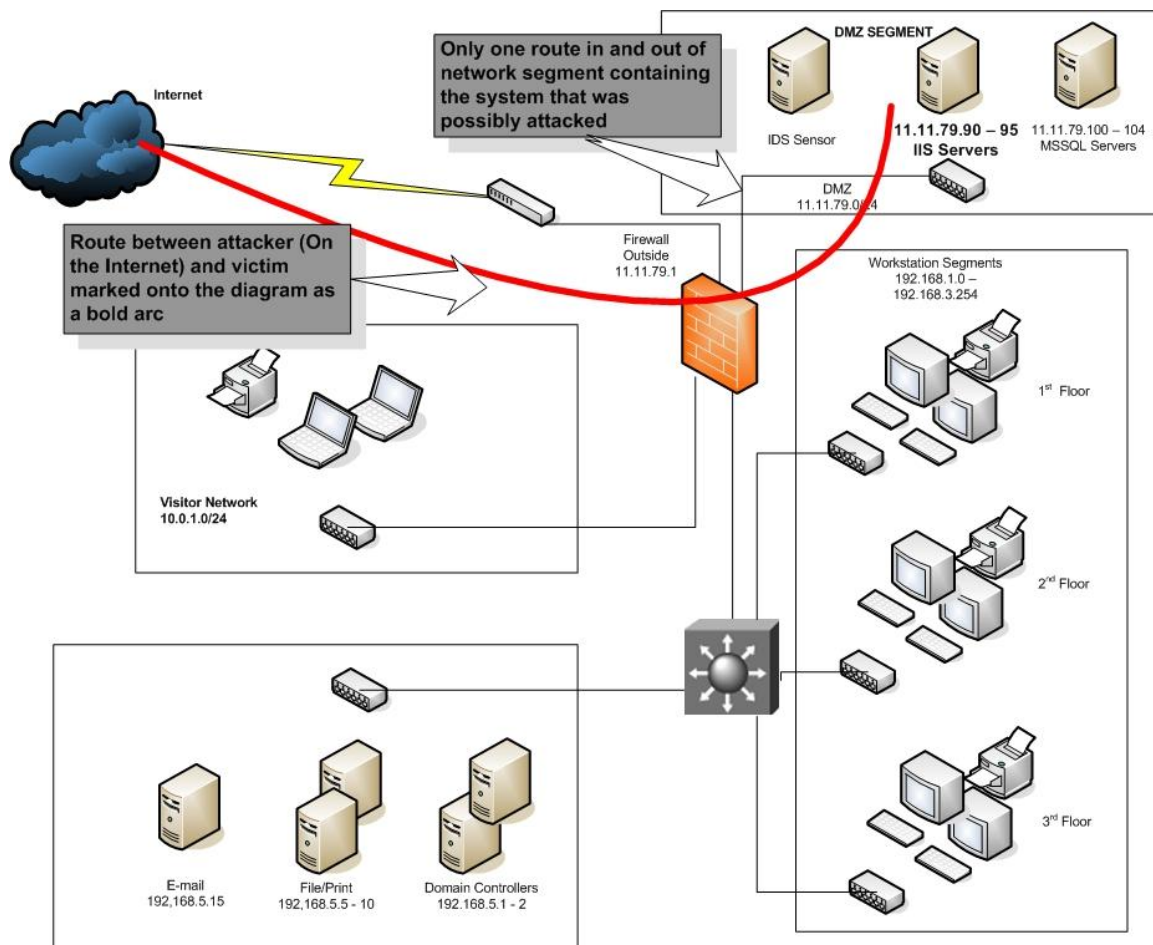
You can use the current observations from the investigation to add source and destination systems to your profile of activity. This is illustrated below, following the example from the previous topic.



Mapping Observed Activity to Network Traffic Flow, continued

Example: Mapping Traffic Routes

Use a network diagram and the IP addresses of involved devices to map relevant routes. The basic example below uses the IP addresses added to the profile of observed activity on the previous page. In the example, a bold line was used to mark the route between the victim system and the attacker who is assumed to be on the Internet. The only point of ingress/egress from the involved network segment was also marked.



Predicting Artifact Location

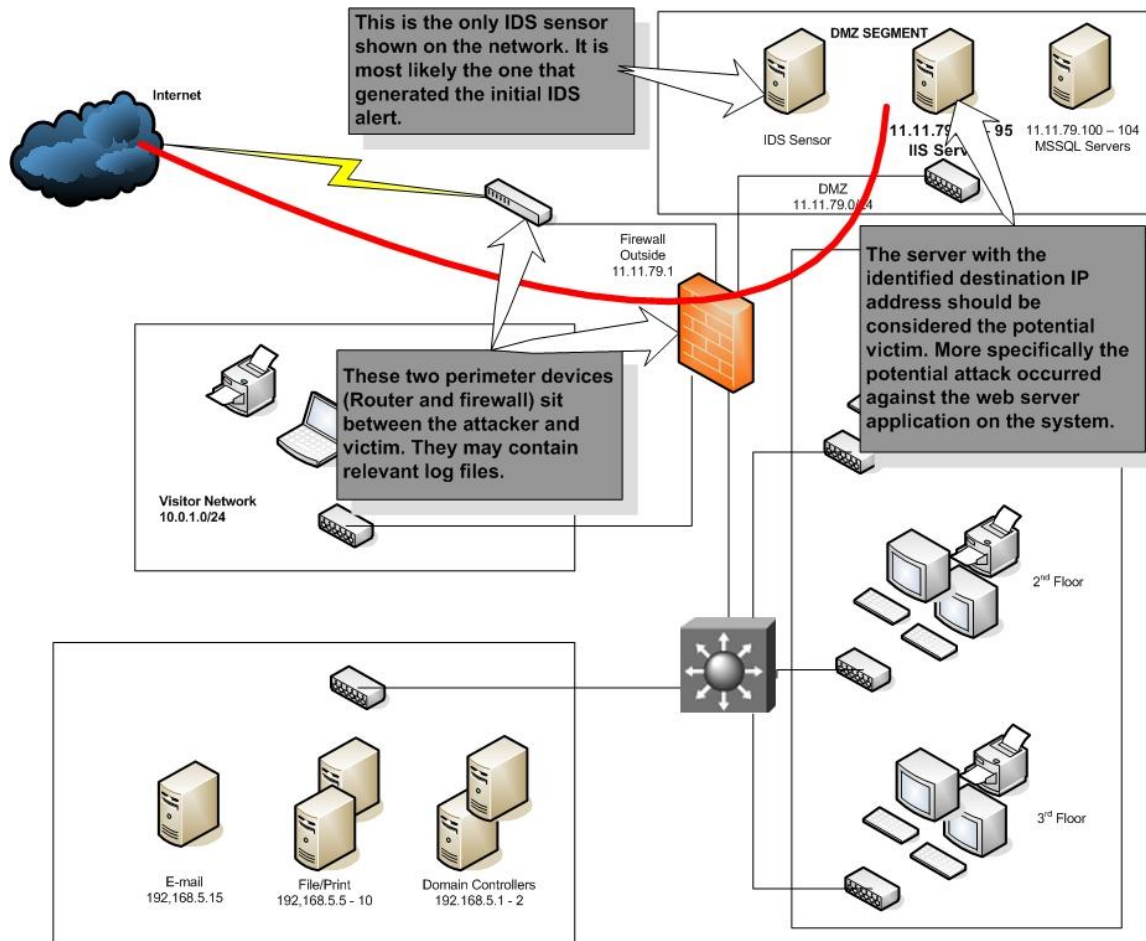
Predicting Artifact Location: Devices Answer the following questions to determine which devices and/or media may contain data regarding any related applications you have identified:

- On what host system is the application located?
- Does the application use local or remote (SAN, NAS, etc.) storage? If remote, identify the associated storage devices.
- Is the host system backed up on a regular basis? Is data backed up to local media, or to a remote system? If local, where are the tapes or other backup media stored after use? If remote, identify the remote backup server.
- Is the application part of a distributed application system (such as a Web server with a database backend)? If so, identify the other applications in the system, and the hosts on which they reside.
- Are there multiple systems that host this application as part of a load-balancing configuration? If so, identify all systems hosting copies of this application.
- Is the application configured to use a proxy device when communicating on a network? If so, identify all associated proxy devices.

Predicting Artifact Location, continued

Predicting Artifact Location: Devices Example

Based upon the map of traffic flow and profile of observed events created previously, devices can be selected which will most likely contain relevant artifacts. The diagram below is provided as a basic example, continuing with the scenario from the previous pages.



Predicting Artifact Location, continued

Predicting Artifact Location: Files and Directories

Answer the following questions to determine which files and directories may contain data about the related application:

- Does the application keep logs? If so, what is the full path to the log storage location?
- Is the application or host system configured to send logs to a remote repository? If so, identify that system.
- What is the full path and name of the files in which the application stores configuration information?
- What is the name and full path of the files in which the application stores persistent and temporary data?
- Does the application require authentication? If so, does it use its own authentication mechanism, or does it forward authentication data to an outside application (such as Active Directory)?

The answers to these questions will vary between different applications and operating systems. Research is required.

Predicting Artifact Location: Files and Directories Example

Following the same example scenario, the potential victim in the attack could be an IIS Web server that contains certain log files to be analyzed. Here are examples of log files that could be extracted from this system:

- Windows Event Logs: The Windows operating system logs (IIS runs on Windows), typically found in C:\[winnt or windows]\system32\config. They will have a “.evt” extension on recent server versions of the Windows OS.
- IIS logs: The log files for the web server application, typically located at c:\[winnt or windows]\system32\logfiles\w3svc1, and will usually have a name of ex*.log.
- Dr. Watson log: A debugging log created by the Windows OS after some program malfunctions, named drwtsn32.log will sometimes contain data if a process was crashed when it was attacked.

Following this example, log files would have to be collected from the other devices identified on the diagram (IDS, firewall, etc.).

This page intentionally left blank.

Lesson 5 – Using Log Analysis to Evaluate an Intrusion Hypothesis

Introduction Once you determine the most likely locations for related artifacts, you can use the techniques presented in this course to analyze collected evidence and evaluate the hypothesis.

Purpose of this Lesson The purpose of this lesson is to describe how log analysis techniques are used to evaluate an intrusion hypothesis.

Objectives After completing this lesson, you will be able to:

- Determine the format of log files
- Use search, filter, and extraction techniques to evaluate a hypothesis
- Record findings and keep track of new leads

In this Lesson The following table shows the contents of this lesson.

Topic	See Page
Hypothesis Evaluation	9-38
Acquiring Target Log Files	9-39
Reviewing Target Log Formats	9-40
Establishing Search/Extraction Criteria	9-41
Searching Target Logs and Extracting Relevant Data	9-42
Recording and Correlating Findings	9-43
Keeping Track of New Leads	9-45

Hypothesis Evaluation

Hypothesis Evaluation

A hypothesis is evaluated using digital forensic data acquisition and analysis techniques. A general process for this is listed below, and will be described in further detail in the following pages.

1. Acquire target log files
2. Review the format of collected logs
3. Establish search/extraction criteria based upon predicted artifacts and log format
4. Search log files and extract relevant data
5. Record and correlate findings
6. Document unexpected findings related to the case (“leads”)

Procedure Selection

There are many technical procedures you can implement to accomplish each of the tasks listed above. For instance there are multiple methods of searching and filtering log files presented in this text. Applied procedures should meet the following criteria:

- The procedure can locate specific potential artifacts identified in the previous step in the Scientific Method.
- The procedure should have been tested and peer reviewed
- The procedure should be repeatable
- The procedure should be as objective as possible

Acquiring Target Log Files

Acquiring Log Files

You will most likely acquire log files through one of the following methods:

- Log files may be provided directly to you by an incident responder or network administrator who collected them from the original source media.
- You may obtain a physical or logical image of the original storage media containing the log files, and then extract the logs from that image.
- You may logically copy log files from the source system or device using a live forensic tool such as LiveWire Investigator.

Regardless of the method, ensure that the logs are collected in a sound manner in accordance with maintaining the integrity of the evidence to the best of your ability. This includes hashing any log files you receive and comparing those hashes against those provided by the responder.

Previewing Log Formats

Previewing Log Formats

Before analyzing collected logs, you should first preview the format of those logs to ensure that you know how to read them properly and use the correct methods for searching them. Search techniques are significantly different between text and binary logs. Furthermore, information may be presented in different forms in text logs. There may be different field formats and field and record separators. For example, time may be represented in a 24-hour format, or in a 12-hour format with AM or PM specified.

Determining File Type

The first step in previewing log format is to determine the file type. This can be done by:

- Viewing the file extension, and correlating it to a file type. For instance a file with a “.EVT” extension is a Windows Event Log file and should be viewed in the Microsoft Event Viewer application. If unfamiliar with the extension, research it online.
- Using the GNU “file” command on a Linux, Unix, or OS X based computer. This does not rely on file extensions, so it is most useful if the log file does not include an extension. Usage example:

```
[prompt]# file evidence.log
```

- Attempting to view the file with a text log viewer, such as notepad.exe, or a GNU command such as “cat”, “more”, “tail”, etc. If successful you will have determined that the file is a text log and viewable as such.
- The File – Open dialog in Wireshark will display the format of binary capture files that it recognizes when you highlight the file.

Determining Data Format within a Log

Once you know the file type for each log, you should identify the format of the data within. For network traffic capture logs, this is relatively uniform. Text logs will vary. When presented with a text log, you should:

- Determine if the records include one line or multiple lines.
- Identify the field and record separators.
- Determine where common data types (IP addresses, port numbers, date/time, etc.) are located in each record, if anywhere. Also determine if these locations are always the same, or variable.

Establishing Search/Extraction Criteria

Search/Extraction Criteria

Your general goal will be to search for and extract log entries, or portions of log entries that support or contradict your hypothesis. To do this, you find entries that include either observed or predicted artifacts. Typical criteria for finding these entries include:

- The known or estimated time frame
- Observed or predicted source or destination IP addresses
- Messages that correlate with observed or predicted activity
- Observed user name or alias
- Observed or predicted network protocols or traffic types
- Any combination of the criteria mentioned in the items above

These examples are guides. There is no hard and fast rule for choosing specific search and extraction criteria. Rather the criteria must be selected based upon the potential for resulting data to further the investigation of the hypothesis.

Searching Log Files and Extracting Relevant Data

Searching Log Files and Extracting Data

At this point, you will use the search and extraction techniques presented earlier to obtain the target data. When performing these tasks, maintain focus on the search/extraction criteria you have established. Investigative tangents based on your intuition can be beneficial, but should be kept to a minimum when attempting to process large amounts of data.

Recording and Correlating Findings

Recording Findings

As noted earlier, your method for recording findings will vary depending upon the approved data repositories used by your organization. This course will use a spreadsheet document provided by your instructor.

Correlation: Timeline Unification

The main task of correlation is the establishment of a unified timeline. This is accomplished by one of the these methods:

- 1) Normalizing all log files to a synchronized time and time notation. This involves changing the actual date/time stamps in a log file and should only be performed on working copies of a log. Investigators must ensure that they do not perform this step on original evidence as the information will be permanently altered, which could prevent it from being admissible in court.
- 2) Record all events identified during initial observation and subsequent testing/evaluation into a single timeline, adjusting the time on each event as necessary as it is recorded.

If you perform these tasks, the resulting timeline can be read sequentially, providing a top-level view of events from all sources. The first option should only be used if you have access to a tool that can reliably interpret and skew the date/time stamps from all log formats, or if you can script this action yourself. Otherwise, use the second option.

Note: Sawmill can skew log file date/time stamps in 1 hour increments.

Recording and Correlating Findings, continued

Correlation: Event Verification Verify events by checking each log entry for another recording of the same event from other sources. Some useful techniques for correlation include:

- Any network event can be correlated with the data in a full network traffic capture, if it is available.
- Verify Web browser history with proxy server logs. Both will record URL access and the associated times.
- Verify IDS scan alerts with firewall logs. A firewall will often block and log some of the scan packets.
- Verify IDS password attacks alerts with authentication logs, such as the Windows Security Event Log or /var/log/secure, etc.
- Verify e-mail header date/time stamps with e-mail gateway or e-mail server logs.

Correlation: Using Event Verification to Synchronize Times The dates/times for events verified against multiple sources can also be compared to see if there is a time skew between the data sources. For instance, the IE history for a user may show access to Web mail occurring at 2105, but the Web proxy shows that access occurring at 2135. The analyst could use these two events to determine that the proxy server clock was most likely set 30 minutes behind the clock on the subject system.

Ideally, there would be multiple events verification that could be used to confirm time skew. More verified events produced increased confidence in the time skew established by the investigator. Additional methods for date/time correlation include:

- Compare date/time stamps embedded in files with the file system date/time stamps.
- Compare the date/time stamp of the last entry in a log with the file system last accessed time.
- If an event involved access to one or more files, check the appropriate file system date/time stamps for that file. For instance, if an IDS alert indicated the traversal of a Windows command shell banner over the network, check the NTFS/FAT last accessed time for cmd.exe, or for the prefetch file corresponding to cmd.exe.

Keeping Track of New Leads

Unexpected Findings

You will often discover information that was not directly predicted during the initial analysis of your hypothesis. This information can be called a “lead.” This information is sometimes related to your hypothesis, and other times be important, but outside the current path of your investigation. In either case, such information should be recorded so that it can be followed up as needed.

Lead Tracking

New leads should be documented. In addition to your investigative notes, leads should be recorded in your Attribute List spreadsheet along with other relevant data. However entries representing leads should be marked as to whether or not they are relevant to a current working hypothesis.

Marking these entries makes it easier for you to review your results at a later time to determine if you need to modify your current hypothesis or create a new one. Methods for annotating an attribute entry as a lead include:

- Highlighting the entry in a different color
- Listing lead entries on a separate page, tab, table, etc.
- Using a column in a table to mark entries as leads

This page intentionally left blank.

Module 10

Log Sources

Overview

Knowing where the logs of interest reside on a system is a key piece of information when starting a network investigation. In this module, you will see some of the typical locations of logs for select applications and systems.

Purpose of this Module

You will learn where to look in Windows, Linux, Solaris and general IDS systems for logs of interest. In some cases, you will look at the typical contents of these logs to give you a better understanding of them.

Objectives

After completing this module, you will be able to:

- Describe the storage locations of typical log files
- Discuss some of the log file formats
- Recognize IDS logs and their contents

In this Module

The following table shows the contents of this module.

Topic	See Page
Lesson 1 – Windows Log Sources	10-3
Lesson 2 – Linux Log Sources	10-9
Lesson 3 – Solaris Log Sources	10-13
Lesson 4 – Log Searching	10-15
Lesson 5 – IDS Logs	10-19

This page intentionally left blank.

Lesson 1 – Windows Log Sources

Introduction This lesson will cover the most common logs found in a Windows environment.

Purpose of this Lesson The purpose of this lesson is to list the locations of log files for the operating system and several standard applications in a Windows environment.

Objectives After completing this lesson, you will be able to:

- Explain where Windows Logs are stored
- Recognize naming conventions of log files
- Identify some of the file formats for these files

In this Lesson The following table shows the contents of this lesson.

Topic	See Page
Windows Logs	10-4
Windows Services Logs	10-6

Windows Logs

Mail

Windows comes with a version of Outlook or Outlook Express mail client. The default location for these log files in Windows 2000, Server 2003 and XP is inside each user's profile.

The log concerning Outlook's MAPI accounts, typically used by most users, is found at: C:\Documents and Settings\username\Local Settings\Temp\Opmlog.log

If a user has established a Hotmail account in Outlook, these events will be logged in:

C:\Documents and Settings\username\Local Settings\Temp\Outlook Logging\Hotmail\http0.log.

Microsoft SQL Databases

Microsoft SQL Server, one of the most popular database services used in businesses, stores its log files in the directory:

C:\MSSQL\LOG

In this directory, you will find the following log files:

- **ERRORLOG** – MS SQL's default error log file. If logging is configured to create new files on a routine basis, or if the file grows too large, additional error logs will be created with a sequential number appended to the end.
- **SQLAGENT.OUT** – Can contain information generated by the SQL programmer or messages generated by default in the administrative panel. These files can have version numbers at the end as well. The file with the OUT suffix is the current log.
- **SQLDump9999.txt** and **SQLDump9999.mdmp** – These are special dump files that can be generated if the SQL Server crashes or terminates unexpectedly. The information in these files generally contains memory and data pointers at the time of the failure. It is also possible for the administrator, or attacker to force the generation of these files under special circumstances.

Note: If the server is running, you may not be able to copy or open the current log files.

Windows Logs, continued

MySQL

MySQL is free, open source database application that is also popular on many Windows systems. The default location for installation of MySQL in Windows is:

C:\Program Files\MySQL\MySQL Server X.X

In this folder name, X.X is the software's version number. Under this folder are the following directories and logs:

- *bin*: Contains the client programs and server program
- *data*: Holds the log files and the actual databases
- *share*: Has the error message files

The error filename will typically start with the network host name of the system MySQL is running on and end with the .err suffix. For example, enron.err.

Microsoft Access

Microsoft stores any errors generated by Access in the Windows Event logs. For information on how to retrieve these logs, refer to the System Logs section below.

Windows Services Logs

Internet Information Server (IIS)

IIS is a service used by millions of Windows-based servers to host Web, FTP, and e-mail services. Depending on the version of IIS in use, these logs can be found in different locations. While IIS will normally store its logs into a default folder, this location can be easily changed in the administration control panel.

For IIS versions 4 and 5, found on Windows NT 4.0 and Windows 2000, log files will be stored in: C:\winnt\system32\logfiles

For IIS version 6 and 7, found on Windows XP and newer systems, log files will be stored in: C:\windows\system32\logfiles

Log file names will be named “W3SVC” followed by the Site Instance ID, which is numbered sequentially for each service. For example, the first web site log files will start with W3SVC1, and the second will be W3SVC2.

Because all Web enabled services originate in the IIS service, FTP and DNS messages will be mingled in this same file if their services are active.

Windows Services Logs, continued

System Logs

Almost all other services that originate in the Windows environment will log entries into one or all of the standard Event Logs for the system.

These logs are divided into the Application, Security, and System logs. If you want to use or view these logs, you must use the Event Viewer that is available from the Administrative Tools Control Panel.

Unless you use a specialized tool like the Event Viewer, the native log files are stored in a mixed binary format, making standard text based tools ineffective. You can choose the log file you are interested in from the menu in Event Viewer and then choose Export List from the Actions drop down menu to export the log as:

- Tab Delimited text
- Comma Delimited text
- Tab Delimited Unicode text
- Comma Delimited Unicode text

Once exported, these files can be filtered and searching using tools like Grep and Findstr.

Directory Services

If Directory Services is configured for diagnostic logging events that generate a log event, these events will be found in the Event Viewer with the System, Application and Security logs in a separate table called Directory Services.

Remote Logs

Looking at mounted share locations and names may give you an indication that logs are being stored remotely. Examination of these remote shares may give you folder and file names which will indicate what types of logs are being stored remotely.

This page intentionally left blank.

Lesson 2 – Linux Log Sources

Introduction This lesson will cover the common and most used logs found in a Linux environment.

Purpose of this Lesson The purpose of this lesson is to list the locations of log files for the operating system and several standard applications in a Linux environment.

Objectives After completing this lesson, you will be able to:

- Explain where Linux logs are stored
- Identify naming conventions of log files
- Recognize some of the file formats for these files

In this Lesson The following table shows the contents of this lesson.

Topic	See Page
Linux Logs	10-10

Linux Logs

Mail Logs

Because Linux is based on the Unix style kernel, mail services are provided by sendmail processes. Logs for these services can usually be found in the file:

`/var/log/maillog`

Database

MySQL is the most popular database program within the Linux community. MySQL logs can typically be found in the `/var/log/mysqld.log` file.

Services

Linux services are usually found in the following directories and files:

- `/var/log/message`: General messages and system related errors
- `/var/log/auth.log`: Remote Login Authentication logs
- `/var/log/secure`: Remote Login Authentication log
- `/var/log/kern.log`: Kernel logs
- `/var/log/cron.log`: Crond logs, for services that start automatically
- `/var/log/httpd/`: Apache web server access and error logs directory
- `/var/log/boot.log` : System boot log
- `/var/log/utmp` or `/var/log/wtmp` : Binary Login history file
- `/var/log/yum.log`: Yum log files to track installed and uninstalled applications

Directory Management

Linux doesn't support Microsoft Active Directory directly, but there are many third party add-on tools available providing this service. You will need to seek out documentation for the specific AD tool and determine the location of logs for each tool.

Linux Logs, continued

System Logs

Most Linux system log entries are located in the `/var/log/message` file.

Remote Logs

Looking at mounted share locations and names can give you an indication that logs are being stored remotely. Examination of these remote shares can give you folder and file names which will indicate what types of logs are being stored remotely.

This page intentionally left blank.

Lesson 3 – Solaris Log Sources

- Introduction** This lesson covers the common and most used logs found in a Solaris environment.
- Purpose of this Lesson** The purpose of this lesson is to list the locations of log files for the operating system and several standard applications in a Solaris environment.
- Objectives** After completing this lesson, you will be able to:
- Explain now where Solaris logs are stored
 - Recognize naming conventions of log files
 - Identify some of the file formats for these files
- In this Lesson** The following table shows the contents of this lesson.

Topic	See Page
Solaris Logs	10-14

Solaris Logs

Mail

Depending on the version of Solaris you are examining, you may find a file in the /etc directory called syslog.conf, and it may have the location of sendmail logs listed inside.

Many ISP's have gone to custom mail software and you may have to seek documentation on the software to determine the log file location.

Databases

If MySQL is installed on the Solaris system you will find the logs in the default locations of either /usr/local/mysql/data or /opt/mysql/mysql/data.

Oracle is a popular database for Solaris systems. You will need to determine the version and release level of Oracle software and then search for the default installation location of log files.

Services

In Solaris, most services put log messages in the /var/adm/messages log file. It is the general catch all file for log entries in a Solaris environment.

Directory Management

Solaris doesn't natively support Microsoft Active Directory directly, but there are numerous third party add-on tools available providing this service. You will need to seek out documentation for the specific AD tool and determine the location of logs for each tool.

System

Traditionally all system log files will be located in the /var directory in a Solaris environment. You will not be able to open files that are in use. There usually are several nested directories of log files under the /var directory and your investigation may show that some or all of these files may have information of evidentiary value.

Remote Logs

You may have to search for pipes and hard links to mounted volumes to discover whether logs are being stored remotely on a Solaris system. In this environment, you may want to locate a certified Solaris administrator to discover some of the obfuscated links.

Lesson 4 – Log Searching

Introduction This lesson presents several ways to manually search through a log file.

Purpose of this Lesson The purpose of this lesson is to show ways in which you can use commonly available tools to search log files.

Objectives After completing this lesson, you will be able to:

- Explain how to use the findstr command
- Describe how to use Grep/Egrep
- Explain the basics of regular expressions

In this Lesson The following table shows the contents of this lesson.

Topic	See Page
Log Searching	10-16
Regular Expressions	10-17
Regular Expressions: Literal Characters	10-18

Log Searching

Overview

Flexibility is the most important feature for log file searching tools. You will encounter a wide variety of log files that will require you to search for different types of values. Your tools and techniques must be usable regardless of the log type and the value for which you are searching.

GREP

One of the primary applications used for searching and filtering text logs is GREP (Global Regular Expressions Print), and its newer version EGREP (“Extended”GREP). These applications use regular expressions to define search parameters. These applications are used because regular expressions are the most common method for defining search parameters and they are used in many other popular applications, such as PERL, Snort, and EnCase.

While GREP is typically found in Unix, Linux, and OS X environments, there are versions available for the Windows operating systems.

FINDSTR

While Windows does not natively ship with GREP, it does include a similar command line utility called Findstr that can help find specific strings of text in a log file or other type of text file. Typing “findstr /?” at the command prompt will display the quick help screen of options and the command format.

Options that may be of interest are:

- /I to disable case sensitivity
- /S to search all files in the current directory and subdirectories
- /R to allow the use of regular expressions
- /N to print line numbers
- /G:filename to use a file of key strings to search for

Regular Expressions

Introduction

Regular expressions are patterns used for executing searches and filters. This is accomplished by combining literal text and special characters, called *metacharacters*, to create a pattern used to search files.

Examples of items that have a set pattern and can be identified with regular expressions include:

- IP addresses
- Dates and time
- Phone numbers
- URLs
- Credit card numbers
- Social Security numbers

For example, an investigation may require that all IP addresses be extracted from a set of logs and put together in a central list. You would not search for a specific IP address because that search would miss IP addresses with different values. The search would have to be for any number that matches the decimal representation of an IP address, which consists of 4 numbers, 1-255, that are separated from each other by periods. Regular expressions can be used to accomplish this and other similar tasks.

You need to understand regular expressions and how to use them effectively to search or filter the wide range of text logs. Many tools incorporate regular expression engines into their standard functionality. Some GNU command line tools, such as grep/egrep, sed and awk, as well as many text editors, allow searching and/or replacement of text through the use of regular expressions. Regular Expression engines and syntax may vary slightly from product to product.

The basic syntax for egrep is:

```
[prompt]# egrep "<expression>" <target log file>
```

Regular Expressions: Literal Characters

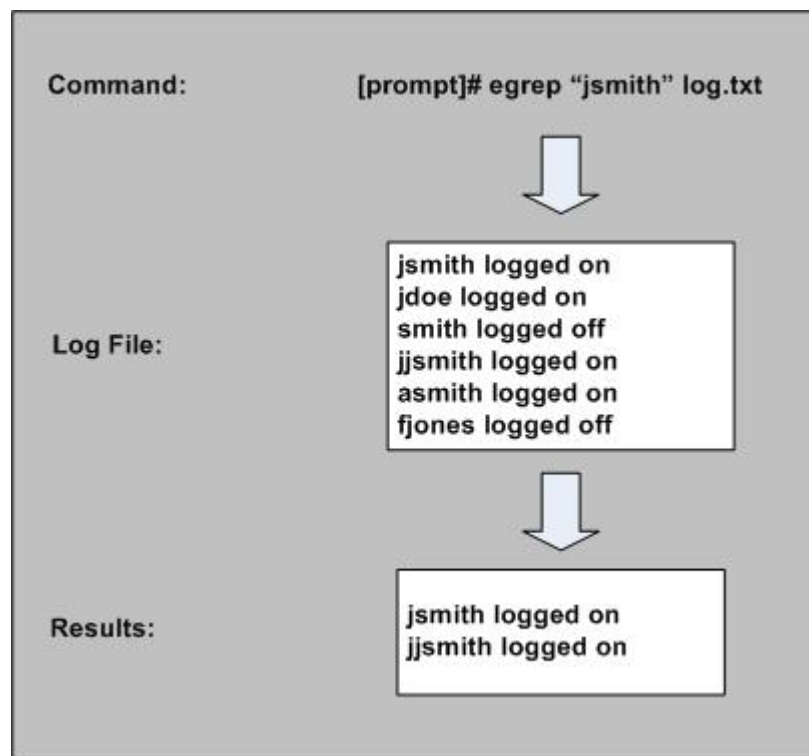
Literal Character Searches

The simplest type of regular expression is the literal representation of the target value. For example, you could search for the word “jsmith” in the file log.txt by simply telling egrep to search for the string “jsmith”:

```
[prompt]# egrep "jsmith" log.txt
```

Most programs search a file one line at a time. This means that the command line above will return each line in a file that contains the string “jsmith” to whatever output is specified.

An example of a search for a literal string is shown below. In the example, a log file is searched for the string “jsmith” using egrep, and the results are displayed. Notice that any line that included the string “jsmith” was matched, even the one that begins with “jjsmith.” In a search for a literal string, it does not matter what is before or after the target, only that the target exists within the line. This means that the target “jsmith” could be a stand-alone word, or just part of a word, such as “jsmithsonian.”



Lesson 5 – IDS Logs

Introduction This lesson will cover Intrusion Detection System logs.

Purpose of this Lesson The purpose of this lesson is to increase your understanding of typical IDS logs and to introduce the Snort tool.

Objectives After completing this lesson, you will be able to:

- Recognize the importance of IDS logs
- Explain how Snort is used

In this Lesson The following table shows the contents of this lesson.

Topic	See Page
IDS Logs	10-20

IDS Logs

Intrusion Detection Systems

Intrusion Detection Systems (IDS) have become prolific and are found in many networked environments. You will probably find that most of the logs generated by these systems are binary rather than text files. When you find these files, you may have to use a proprietary program to view or convert the file to text. Some of these IDS will save logs in the libpcap format. This format enables you to use packet sniffer tools like Wireshark to open, view, and export the files as needed.

Snort

Snort is a popular IDS and intrusion reporting tool. Along with being a popular free application, Snort allows administrators to flag alerts on both live traffic and traffic captured with a packet sniffer. After parsing through traffic, Snort will generate a text log displaying all of the alerts of suspicious traffic it encountered. An example of such an alert is displayed below:

```
[**] [1:2001689:5] BLEEDING-EDGE WORM Potential MySQL
bot scanning for SQL server [**]
[Classification: A Network Trojan was detected] [Priority: 1]
08/18-12:21:55.172252 <remote IP>:49812 -> <local IP>:3306
TCP TTL:93 TOS:0x20 ID:256 IpLen:20 DgmLen:40
*****S* Seq: 0xBE7728D2 Ack: 0x0 Win: 0x4000 TcpLen: 20
[Xref => http://isc.sans.org/diary.php?date=2005-01-27]
```

Be aware that Snort requires a complex set of steps to configure it properly, and this configuration will change with each type of log or capture you feed it. By default, all of Snort's log files on a Linux, Unix, or OS X system will be found in: /var/log/snort

Module 11

Log Analysis

Overview When identified, log data must properly formatted and assembled into reports. Log entries can be used directly as items of evidence, or assembled into other forms of data, such as statistics, charts, graphs, and other representations.

Purpose of this Module You will be introduced to methods for manipulating log data into formats that can be easily analyzed for pertinent information.

Objectives After completing this module, you will be able to:

- Generate statistics from log data
- Format log data into report-friendly formats
- Form visual charts and graphs with log data

In this Module The following table shows the contents of this module.

Topic	See Page
Lesson 1 – Binary Traffic Analysis	11-3
Lesson 2 – Manual Log Analysis	11-23
Lesson 3 – Automated Log Analysis Tools	11-29

This page intentionally left blank.

Lesson 1 – Binary Traffic Analysis

Introduction

Binary logs require different filtering and searching techniques than those that are used with text logs. Due to the size of binary logs and their required processing power, it is often more efficient to filter binary network captures with command line tools.

Purpose of this Lesson

You will learn techniques for filtering and searching binary logs of network traffic using command line tools.

Objectives

After completing this lesson, you will be able to:

- Describe the types of criteria that can be used to filter binary logs
- Convert binary logs to text files
- Demonstrate how to filter and search binary logs with Wireshark

In this Lesson

The following table shows the contents of this lesson.

Topic	See Page
Introduction to Wireshark	11-4
Converting Binary Logs to Text Format	11-5
Filtering and Searching in Wireshark	11-6
Filtering Data during Capture with Wireshark	11-7
Filtering Displayed Data in Wireshark	11-8
Colorizing Data Using Filters in Wireshark	11-14
Searching in Wireshark	11-16
Generating Statistics with Wireshark	11-17
Exporting Data from Wireshark	11-22

Introduction to Wireshark

Wireshark

Wireshark is a powerful, open source protocol analyzer that can be used to view full network traffic capture logs. Wireshark can:

- Open a variety of binary log formats
- Act as a sniffer
- Translate, or decode, known protocols within a binary log to human readable format
- Display highly detailed information on a frame-by-frame basis
- Search through a capture log for frames that match specific criteria
- Automatically reconstruct TCP sessions

Procedure: Importing Logs into Wireshark

These steps show how to import a binary log file into Wireshark for analysis.

Note: This procedure and all others in this lesson use Wireshark 0.99 version. Be aware that new versions are released frequently and can have menu options in different locations.

Step	Action
1	Open Wireshark.
2	Click File on the menu bar and select Open.
3	In the Open Capture File dialog box, browse to the location of the capture file.
4	Left click the file name one time to highlight it.
5	Click Open .

Viewing Binary Logs in Wireshark

Wireshark displays binary logs in a window with three panes which contain the following information:

- Top Pane: Summary of captured frames, including frame number, date and time, source IP, destination IP, protocol and basic description
- Middle Pane: Decoded protocol header information, organized inversely to the order of each protocol within the OSI model
- Bottom Pane: Full frame contents in hexadecimal on the left side with any included clear text displayed on the right

Converting Binary Logs to Text Format

Binary vs. Text Format

The following lessons of this module will provide instructions for searching and filtering logs in binary format. However, it is sometimes more efficient to change binary logs to text format. By doing so, the logs can be manipulated using text log filtering techniques to quickly find target data.

You can change binary logs to text format with tcpdump. The default output of tcpdump is text format. Therefore, it can be used to read a binary capture and redirect the output to a text file instead of the screen. Here is an example:

```
[prompt]# tcpdump -r log.cap > log.txt
```

In the command line above, tcpdump read the file log.cap, and placed a text interpretation of the contents into the text file log.txt. Note that by default tcpdump does not print the full contents of each packet, just some summary data.

This is not always necessary when performing quick searches and filters on network protocol data. It will sometimes be useful to modify the default output format of tcpdump. Options for doing this include:

Command Option	Description
-A	Print the content of each packet in hex, except for the Data Link Layer
-n	Do not convert numbers to names, such as port numbers to service names, or IP addresses to hostnames
-tttt	Print the date as the first field of the packet before the time
-v, -vv, -vvv	Print more verbose output, progressively increasing with more v's

Filtering and Searching with Wireshark

Filtering Binary Logs with Wireshark

Wireshark offers several filtering and searching options:

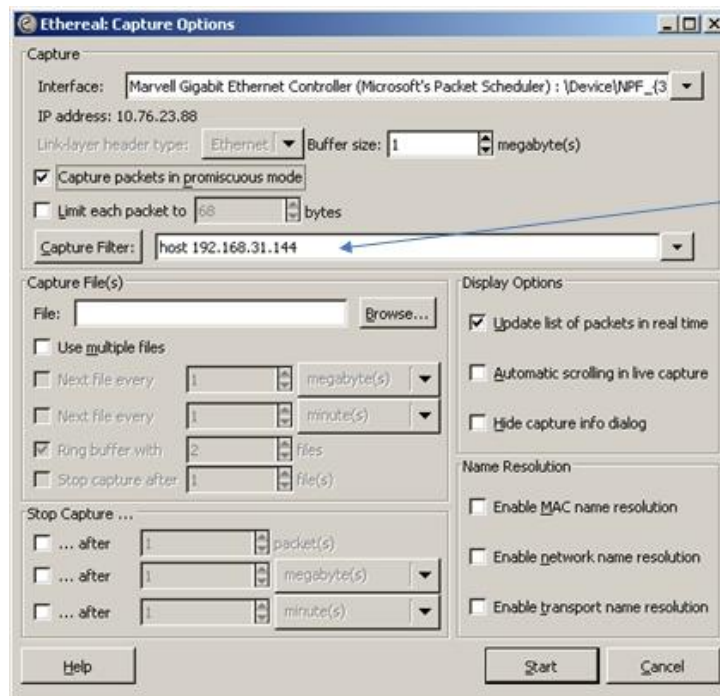
- Capture filters: Interface used to filter data while it is being captured from a network; uses the tcpdump syntax.
- Display filters: Interface used to filter traffic that is currently being displayed by Wireshark.
- Color filters: Interface used to apply colors to certain packets based upon a filter expression.
- Find menu: Standard find menu that allows packet to be searched by hex value or string. Display filters can also be entered here.

Filtering Data during Capture with Wireshark

Procedure: Setting Up a Capture Filter

Wireshark can filter data while it is being captured. This filter menu uses standard tcpdump syntax. Data filtered out through this method never gets stored.

Step	Action
1	Select Capture from the menu bar and select "Options" from the drop-down menu.
2	A window will appear titled "Wireshark: Capture Options." Enter the desired expression into the dialog box next to the "Capture Filter" button. Make sure that the expression is in tcpdump format.
3	Modify other capture options as necessary.
4	Click the Start button to begin capturing.

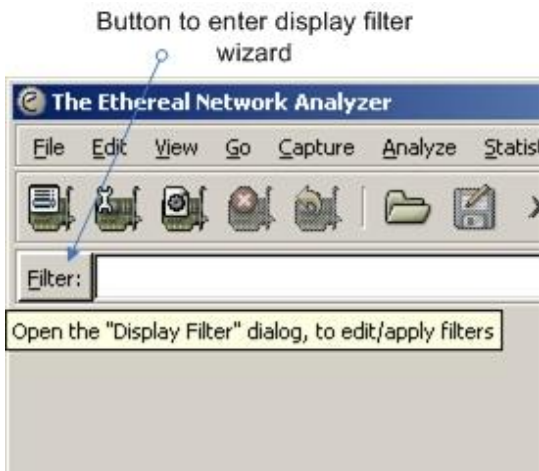
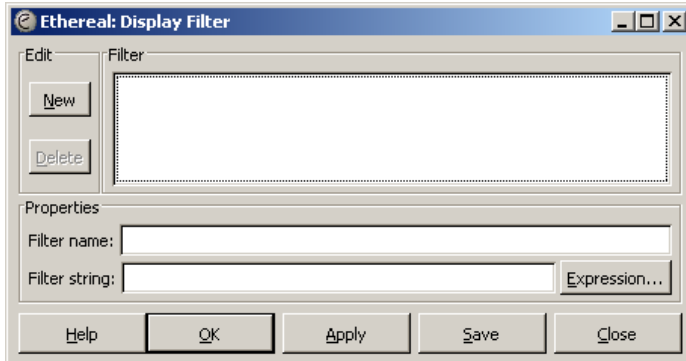


Filtering Displayed Data in Wireshark

Procedure:

Creating a Display Filter in Wireshark

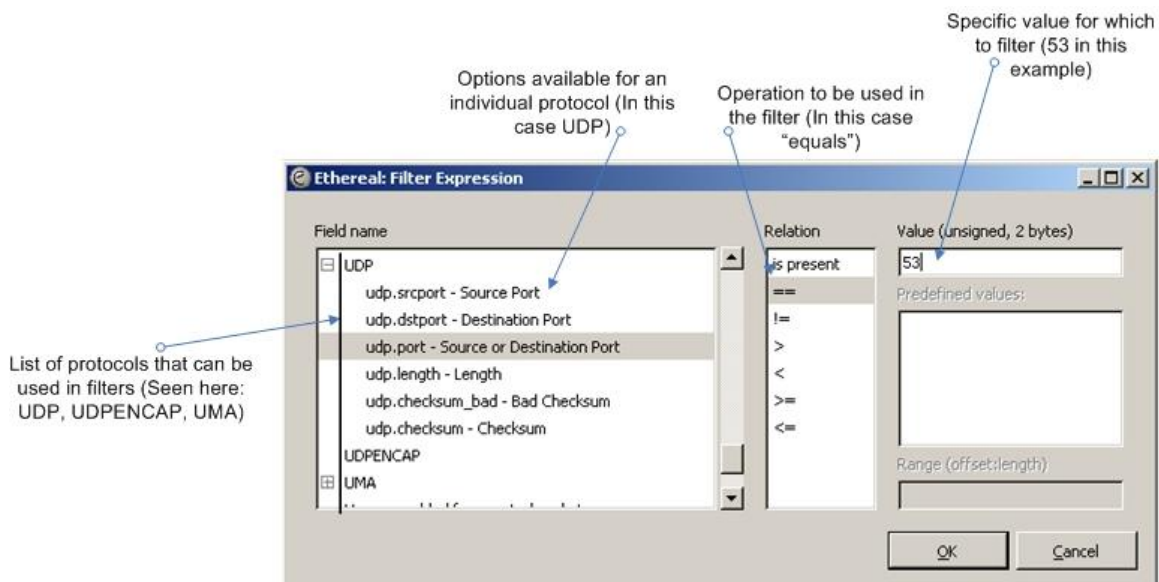
Follow these steps to create and apply a display filter in Wireshark. This filter is used for data that is being displayed in Wireshark. Only the displayed data is changed, not the contents of the log.

Step	Action
1	<p>Click on the Filter button, towards the upper left corner of the Wireshark window.</p> 
2	<p>A new window will appear titled “Wireshark: Display Filter.” Click the Expression button.</p> 

Filtering Displayed Data in Wireshark, continued

Procedure: Creating a Display Filter in Wireshark, continued

Step	Action
3	A window will appear titled “Wireshark: Filter Expression.” Scroll down in the “Field Name” pane until you see the protocol that is targeted for filtering.
4	Left click the arrow beside the protocol name one time to expand the menu of options.
5	Scroll down further and locate the protocol option on which a filter is desired and left click on it. The middle and right panes may change to reflect options available for that protocol option.
6	In the Relation pane of this window, select the desired option by left clicking on it one time.
7	If needed, enter a value into the “Value” dialog box. Click OK .
8	You will be returned to the Display Filter window shown in Step 2. Enter a name for the new filter in the Filter Name dialog box.
9	Click the New button.
10	Click OK to complete the procedure.

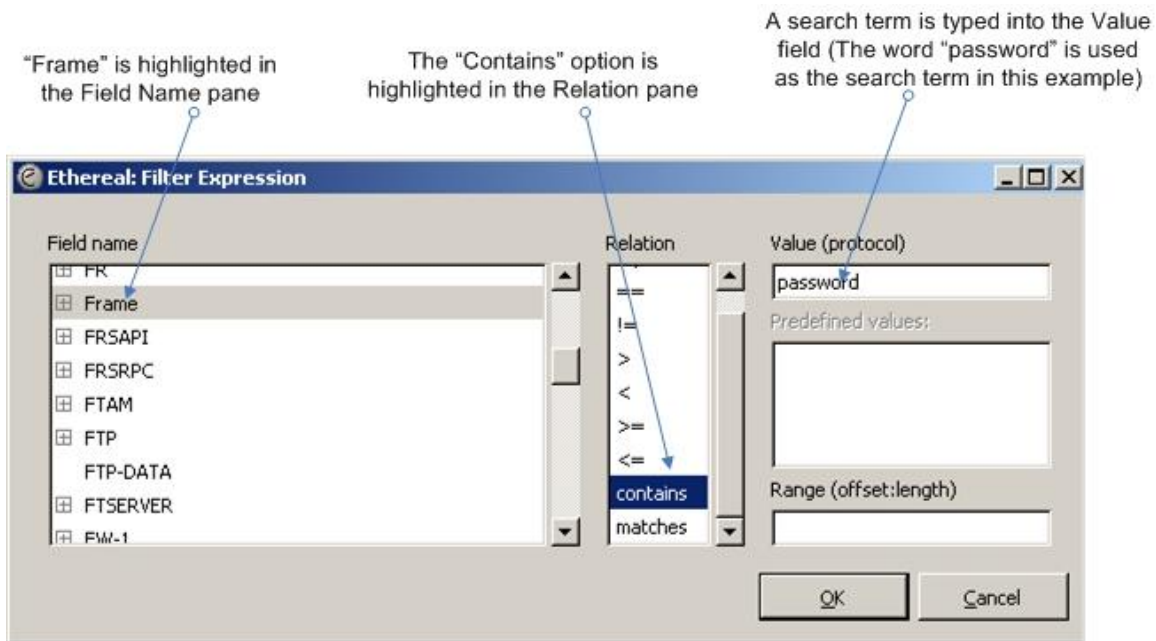


Filtering Displayed Data in Wireshark, continued

Creating a Display Filter for a Keyword

A display filter can be created for a keyword. Unlike a keyword search shown later in this lesson, a keyword filter will change the display so that it only shows packets that contain the search term. This is done with the “frame contains” display filter, which can be used to filter for the presence of a keyword anywhere in a packet.

The “frame contains” filter can be found through the normal display filter wizard. An example of the Filter Expression window from the wizard is shown below. It displays the creation of the “frame contains” filter by choosing “Frame” in the left pane and “contains” in the center pane. The target keyword is placed in the “Value” field.



Creating a Display Filter for a Hex Value

The “frame contains” expression syntax can also be used to filter for hexadecimal values. For instance, the following expression could be used to display packets containing the hex value 0x6d73646f.

frame contains 6d:73:64:6f

The hex value is entered in place of a keyword, with colons used to separate the value into pairs.

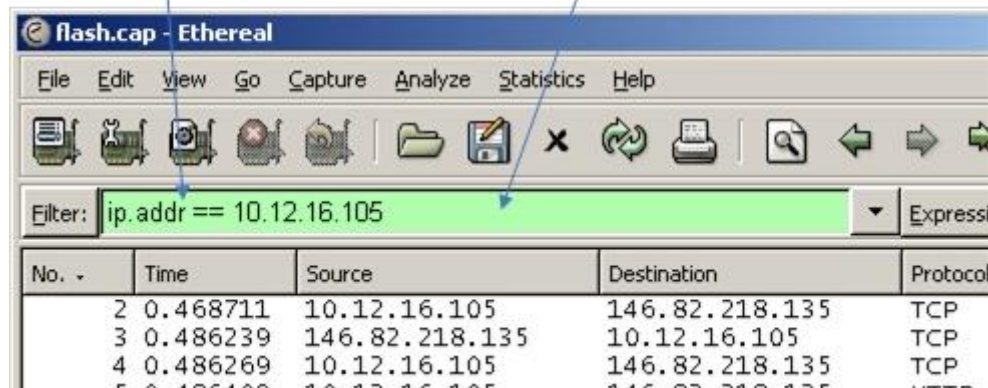
Filtering Displayed Data in Wireshark, continued

Directly Entering Display Filter Expressions

When a filter expression is created using the Display Filter wizard, the text for the filter is entered into the Display Filter field in the main window of Wireshark. Filter expressions can also be entered as text into that field instead of using the wizard. The screen below shows the Display Filter field with a single filter expression displayed. When entering a display filter, if the filter has a valid syntax, the background color of the display filter field will be green, otherwise it will be red.

A filter expression for all traffic to or from the IP address 10.12.16.105

Display Filter field



Filtering Displayed Data in Wireshark, continued

Syntax of Display Filters

Wireshark display filters use a different syntax than tcpdump. The available protocols and filtering options are extensive and cannot all be listed in this text. Here are some common examples.

Operation	Syntax	Example
Source or Destination IP Address	ip.addr == <address>	ip.addr == 192.168.0.1 ip.addr == 20ab:5183:4383::2ff:fee2:7596
Source IP	ip.src == <address>	ip.src == 192.168.0.1 ip.src == 20ab:5183:4383::2ff:fee2:7596
Destination IP	ip.dst == <address>	ip.dst == 192.168.0.1 ip.dst == 20ab:5183:4383::2ff:fee2:7596
Source or Destination Port Number	tcp.port == <number> udp.port == <number>	tcp.port == 80 udp.port == 53
Source Port	tcp.srcport == <number> udp.srcport == <number>	tcp.srcport == 80 udp.srcport == 53
Destination Port	tcp.dstport == <number> udp.dstport == <number>	tcp.dstport == 80 udp.dstport == 53
Protocol	<protocol>	icmp

Filtering Displayed Data in Wireshark, continued

Altering and Combining Expressions

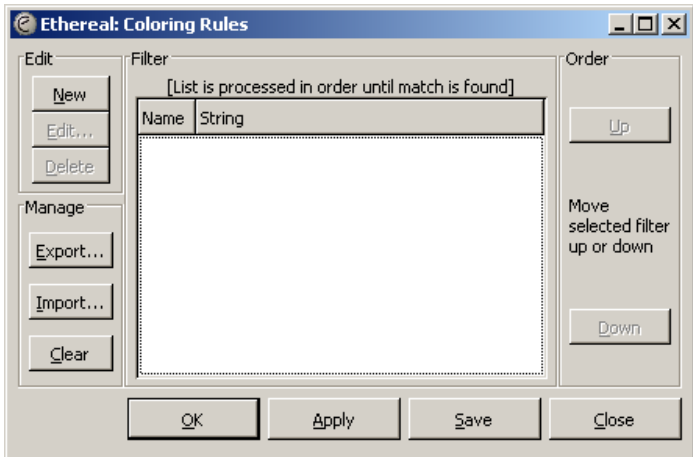
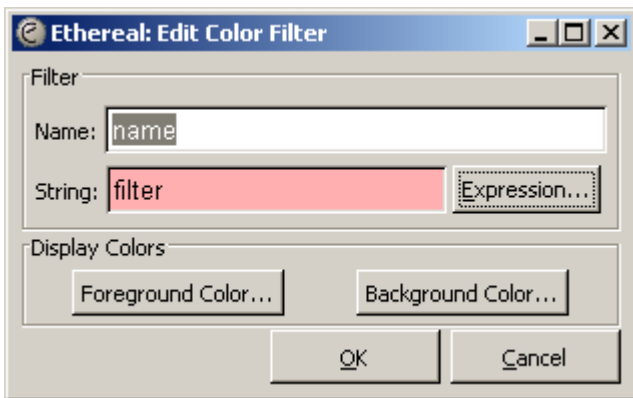
Expressions can be combined in the Display Filter Field and logical operations can be performed on them. Allowed grouping and logic operators include the following.

Operation	Syntax	Example
Combine Two Expressions	<Filter 1> and <Filter 2>	ip.addr == 192.168.0.1 and tcp
Negate an Expression	! <Filter 1>	! ip.addr == 10.0.0.4
Alternate Expressions	<Filter 1> or <Filter 2>	ip.addr == 10.0.0.4 or ip.addr == 10.0.0.5
Compare to Value with "Lesser Than"	<Filter 1> < <Value>	tcp.port < 1024
Compare to Value with "Lesser Than or Equal To"	<Filter 1> <= <Value>	tcp.port <= 1024
Compare to Value with "Greater Than"	<Filter 1> > <Value>	tcp.port > 1024
Compare to Value with "Greater Than or Equal To"	<Filter 1> >= <Value>	tcp.port >= 1024
Compare to Value with "Equal To"	<Filter 1> = <Value>	tcp.port == 1024
Group Expressions with Parenthesis	(<Filter 1> <Operator> <Filter 2>)	tcp.port 80 or (icmp or arp)

Colorizing Data Using Filters in Wireshark

Procedure: Creating a Color Filter with Wireshark

Wireshark can create a filter that does not remove data from the log display, but instead colorizes frames based on the selected criteria.

Step	Action
1	Select View on the menu bar, and then select Coloring Rules.
2	When a new window appears titled Coloring Rules, click New .
	
3	When a new window displays titled Edit Color Filter, click on Expression and the Filtering Expression window displays.
	

Colorizing Data Using Filters in Wireshark, continued

Procedure: Creating a Color Filter with Wireshark, continued

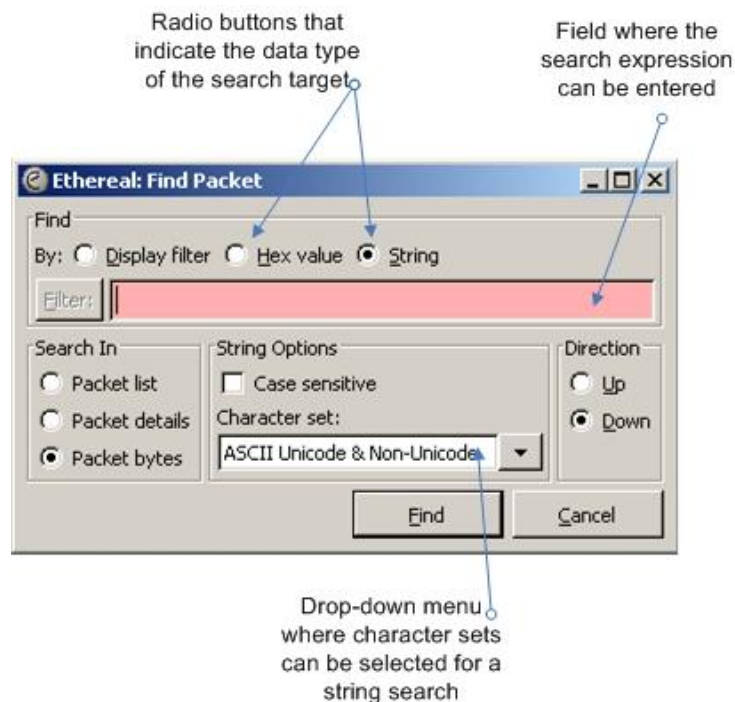
Step	Action
4	This Filtering Expression window is the same one used when creating a display filter (see the previous procedure). Create the necessary filter here in the same way and then click OK .
5	You will be returned to the Edit Color Filter window. Enter a name for the new filter in the Name dialog box.
6	Click Background Color . A new window displays titled “Wireshark: Choose background color.” <div data-bbox="683 795 1378 1337" data-label="Image"> </div>
7	Choose a color by left clicking on the color palette and click OK .
8	You will be returned to the Edit Color Filter window shown in Step 3. Click OK .
9	You will be returned to the Coloring Rules window shown in Step 2. Click OK .

Searching in Wireshark

Procedure: Searching in Wireshark

A standard search for text or hex data can be conducted within Wireshark. This function will not remove frames from the display in Wireshark like display filters do. Instead, the search will scan through the frames and highlight the first frame that matches the search criteria.

Step	Action
1	Select Edit in the menu bar and then select Find Packet from the drop-down menu.
2	A new window will appear titled Wireshark: Find Packet. Select the data type for the search from the radio buttons along the top of the window. Options include: <ul style="list-style-type: none"> • Display filter: Enter a standard display filter • Hex value: Enter a hex value as the search target • String value: Enter a string value as the search target
3	Enter a target value in the field next to the Filter button.
4	Click the Find button and the display will change back to the main Wireshark window. The first matching frame will be highlighted.



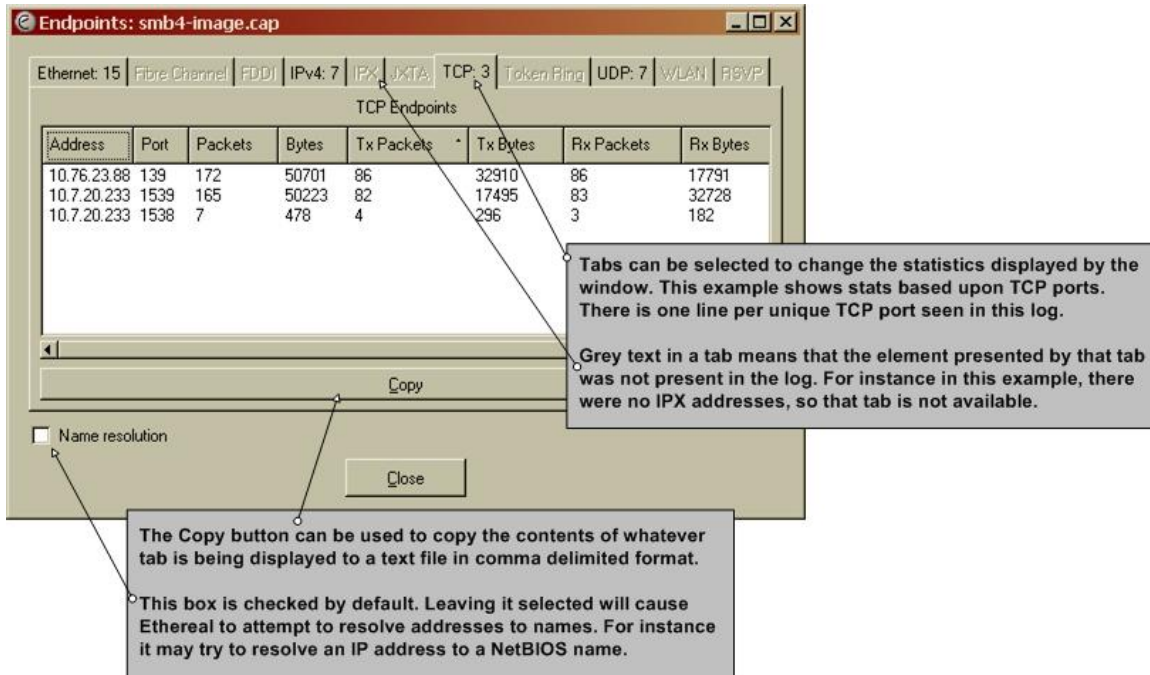
Generating Statistics with Wireshark

Wireshark Statistics Menu

Wireshark has a “Statistics” menu in the menu bar for generating various statistics about log data. Some useful statistics options are described in the following sections.

Endpoints List

The “Endpoints” option in the Statistics menu provides lists of statistics that revolve around addresses and TCP/UDP ports. By selecting the Endpoints option, a separate window will appear to display the statistics. This window is useful to see what IP addresses and ports are seen in a given binary capture. An example is shown below with an explanation of some of the features.



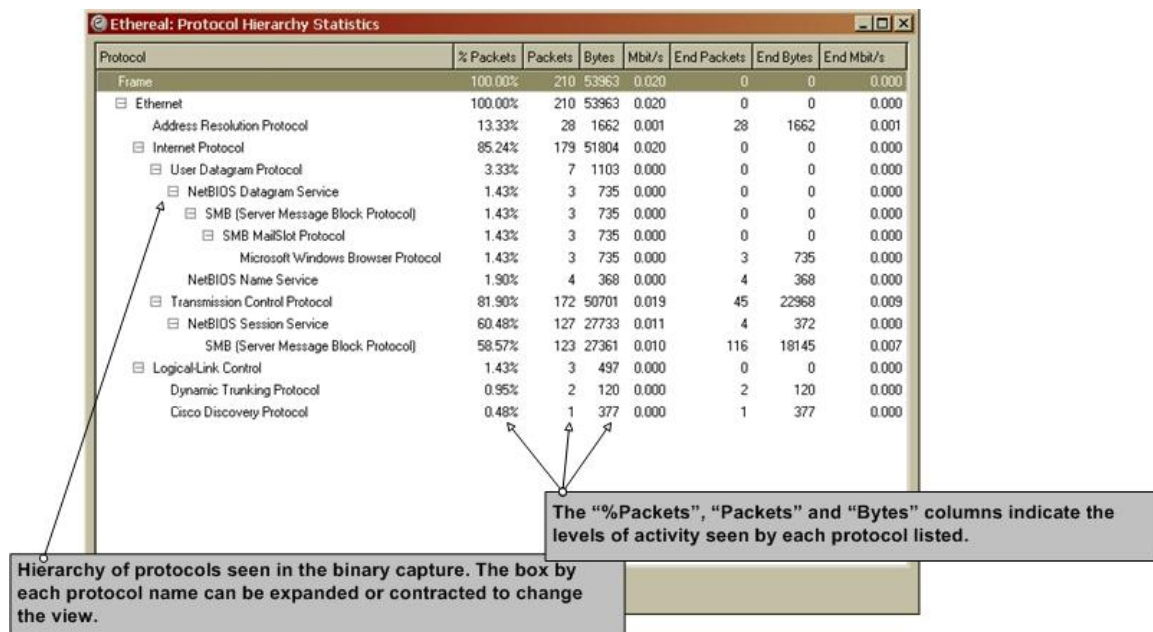
Generating Statistics with Wireshark, continued

Protocol Hierarchy Statistics

The “Protocol Hierarchy Statistics” option in the Statistics menu provides a list of protocols that were seen in a given capture and the volume of protocol activity by each one.

The applications being used on the network can be quickly derived from this information and provide a good snapshot of activity.

However, remember that Wireshark does not recognize all protocols. It may miss a protocol that is used over a non-standard port. An example of this window is shown below.



Generating Statistics with Wireshark, continued

Conversations List The “Conversations” option in the Statistics menu offers lists of source/destination address combinations. Wireshark presents source and destination address combinations that were seen communicating in the capture and the number of packets seen between each pair. Packet volume is even shown for each direction of communication between the pair.

Conversations: flash.cap

Ethernet 3 | File Channel | FDDI | **IPv4: 4** | IPX | JXTA | SCTP | TCP: 7 | Token Ring | UDP | WLAN | NCP | RSVP

IPv4 Conversations

Address A	Address B	Packets	Bytes	Packets A>B	Bytes A>B	Packets A<B	Bytes A<B
10.12.16.105	216.239.39.99	4	234	1	54	3	180
10.12.16.101	10.12.16.105	28	6106	13	2125	15	3981
10.12.16.105	63.236.111.50	33	5446	18	3015	15	2431
10.12.16.105	146.82.218.135	2790	2829604	942	55630	1848	2773974

Copy

☒ Name resolution

Close

Tabs can be selected to change the statistics displayed by the window. This example shows stats based upon IP address source and destination combinations are shown.

Grey text in a tab means that the element presented by that tab was not present in the log. For instance in this example, there were no IPX addresses, so that tab is not available.

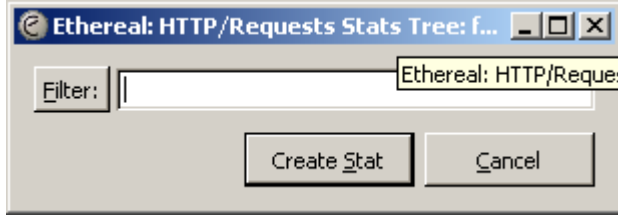
The Copy button can be used to copy the contents of whatever tab is being displayed to a text file in comma delimited format.

This box is checked by default. Leaving it selected will cause Ethereal to attempt to resolve addresses to names. For instance it may try to resolve an IP address to a NetBIOS name.

Generating Statistics with Wireshark, continued

HTTP Requests Stats Tree

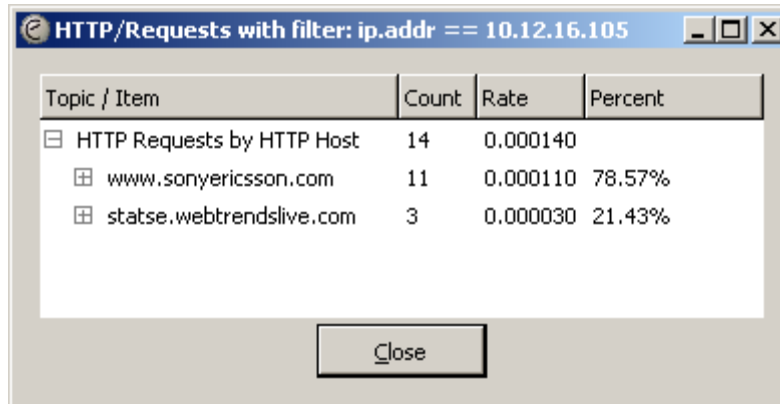
Wireshark can create a custom list of HTTP get requests based upon a specified display filter. For example, if you supply a display filter for a specific IP address, Wireshark shows all get requests for that IP. Creating this statistic requires several steps. Here is the procedure:

Step	Action
1	Select “HTTP” from the “Statistics” menu in the menu bar.
2	When a menu displays, select “Requests...” A new window will appear: 
3	Enter a display filter in the only field. The Filter button can be selected to access Wireshark’s display filter wizard if the desired filter is unknown.
4	Click Create Stat . The statistics window will display.

Generating Statistics with Wireshark, continued

HTTP Requests Statistics Example

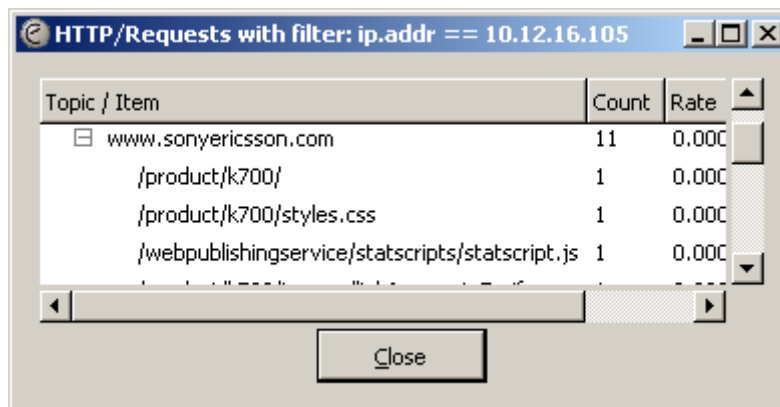
Here is an example of an HTTP Requests Statistics list. The list includes domain names which were found in the traffic. You can also see the percentage of traffic going to each domain name. The example is the result of filtering for all traffic from a single IP address.



Topic / Item	Count	Rate	Percent
HTTP Requests by HTTP Host	14	0.000140	
www.sonyericsson.com	11	0.000110	78.57%
statse.webtrends.live.com	3	0.000030	21.43%

HTTP Requests Stats Tree Example Expanded

You can click the plus sign by each domain name to expand a full list of resources accessed from that domain name. In the following example, the plus sign by `www.sonyericsson.com` was clicked. You can see that multiple URLs at that domain name were accessed by the computer that was the subject of the filter.

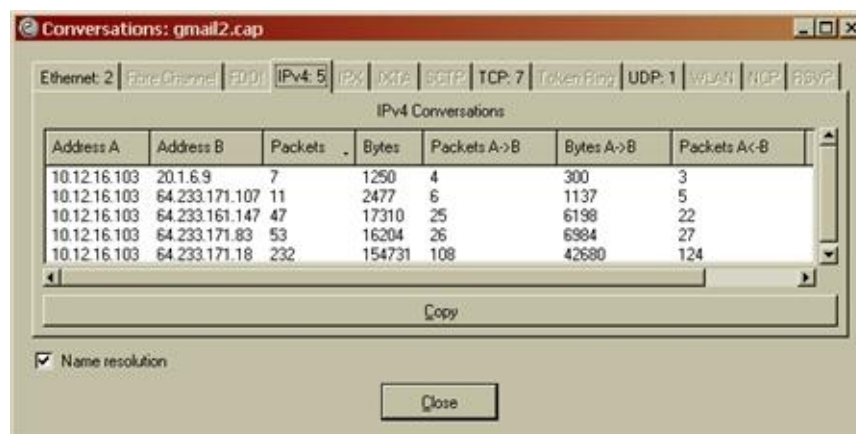


Topic / Item	Count	Rate
www.sonyericsson.com	11	0.000110
/product/k700/	1	0.000010
/product/k700/styles.css	1	0.000010
/webpublishingservice/statscripts/statscript.js	1	0.000010

Exporting Data from Wireshark

Exporting Statistics from Wireshark

You have the option within some Wireshark statistics windows to send a copy of any generated statistics to a file. Statistics windows that are capable of exporting data will have a “Copy” button on the window, as illustrated in the following screen.



Pressing the copy button only puts the data into the copy buffer of the computer with Wireshark. To save the data, it has to be pasted into a file. Use a typical text file application such as Notepad. The data should not be pasted directly into a spreadsheet, because it will all be placed into a single cell.

When pasted into a text file, the data is in comma-delimited format, and includes column headings. However, only the tab currently displayed in the statistic window is actually copied.

For example, the statistic window above is output as:

```
Address A,Address B,Packets,Bytes,Packets A->B,Bytes A->B,Packets A<-B,Bytes A<-B,
10.12.16.103,20.1.6.9,7,1250,4,300,3,950,
10.12.16.103,64.233.171.107,11,2477,6,1137,5,1340,
10.12.16.103,64.233.161.147,47,17310,25,6198,22,11112,
10.12.16.103,64.233.171.83,53,16204,26,6984,27,9220,
10.12.16.103,64.233.171.18,232,154731,108,42680,124,112051,
```

Lesson 2 – Manual Log Analysis

Introduction When automated tools for log analysis are not readily available, this lesson explains how to manually examine and search log files for evidentiary information.

Purpose of this Lesson The purpose of this lesson is to introduce you to ways in which you can search log files manually.

Objectives After successfully completing this lesson, you will be able to:

- Explain how to build keyword lists for searching
- Execute simple searches using EGREP
- Discuss the basic concept of correlation of data

In this Lesson The following table shows the contents of this lesson.

Topic	See Page
Filtering and Searching Text Logs	11-24
Deciding What to Search For	11-25
Example Log	11-26

Filtering and Searching Text Logs

Filtering and Searching Text Logs

Filtering and searching text logs requires the following capabilities:

- Identify all log entries with a specific value or range of values
- Modify the view of one or more log files based upon the existence of an arbitrarily defined parameter

Filtering and Searching Toolset

Flexibility is the most important feature for any tool used to perform filtering and searching. This is because over time you will encounter a wide variety of log files that will require you to search for different types of values. Your tools and techniques must be usable regardless of the log type and the value for which you are searching.

The primary application used in this course for searching and filtering text logs is GREP (Global Regular Expressions Print), and its newer version EGREP (Extended Global Regular Expressions Print). These applications use regular expressions to define search parameters. Regular expressions are the most common method for defining search parameters, and are used in many other popular applications, such as PERL, Snort, and EnCase.

Deciding What to Search For

Keywords

Before searching for data in a log file, you first need to have a clear understanding of what you are searching for. Rarely will a “shotgun” or broad focused search turn up useable data. If you decide on keywords that might be available in the log and would be a possible artifact of the intrusion you are investigating.

Sample Keywords

If you are investigating a person that has attacked a Web server and gained access to the administration area, some of the things you might search for in the Microsoft Internet Information System (IIS) log files might include;

- “Error” or “err”
- “Overflow”
- “Password” or “Pass”
- “Admin”
- “Unauthorized”
- IP addresses of interest

Knowing how Microsoft structures error messages in IIS logs will be a help in deciding the keywords to look for.

WordPad is not Your Friend

A tool like WordPad or Notepad can be used for these types of searches; however, you will find that the data returned is not easily useable and does not allow you to filter the information for clarity or further use.

You can obtain different versions of the Grep command for Windows operating systems from several sites. One GUI version of Grep, WinGrep, is available at <http://www.wingrep.com>.

In the following examples we will be using the Unix version of the Grep command.

Example Log

First Look

In this example, you review logs from a Web server that was exploited. The logs are from the IIS server on the day of the attack.

```
#Software: Microsoft Internet Information Services 6.0
#Version: 1.0
#Date: 2006-05-31 15:51:51
#Fields: date time s-sitename s-ip cs-method cs-uri-stem cs-uri-query s-port cs-username c-ip cs(User-Agent) sc-status sc-su
bstatus sc-win32-status
2006-05-31 15:51:51 W3SVC508294276 10.8.1.39 GET /index.html - 80 - 10.8.5.128 Mozilla/5.0+(Macintosh;+U;+PPC+Mac+OS+X+Mach-
0;+en-US;+rv:1.8.0.3)+Gecko/20060426+Firefox/1.5.0.3 200 0 0
#Software: Microsoft Internet Information Services 6.0
#Version: 1.0
#Date: 2006-05-31 19:58:31
#Fields: date time s-sitename s-ip cs-method cs-uri-stem cs-uri-query s-port cs-username c-ip cs(User-Agent) sc-status sc-su
bstatus sc-win32-status
2006-05-31 19:58:31 W3SVC508294276 10.8.1.39 GET / - 80 - 10.8.1.1 Mozilla/4.0+(compatible;+MSIE+6.0;+Windows+NT+5.1;+SV1;+.
NET+CLR+1.1.4322) 403 14 5
2006-05-31 19:58:42 W3SVC508294276 10.8.1.39 GET /index.html - 80 - 10.8.1.1 Mozilla/4.0+(compatible;+MSIE+6.0;+Windows+NT+5
.1;+SV1;+.NET+CLR+1.1.4322) 200 0 0
2006-05-31 19:59:03 W3SVC508294276 10.8.1.39 GET /index.html/admin.pl - 80 - 10.8.1.1 Mozilla/4.0+(compatible;+MSIE+6.0;+Win
dows+NT+5.1;+SV1;+.NET+CLR+1.1.4322) 404 0 3
2006-05-31 19:59:05 W3SVC508294276 10.8.1.39 GET /index.html/backend.pl - 80 - 10.8.1.1 Mozilla/4.0+(compatible;+MSIE+6.0;+W
indows+NT+5.1;+SV1;+.NET+CLR+1.1.4322) 404 0 3
2006-05-31 19:59:05 W3SVC508294276 10.8.1.39 GET /index.html/backup.pl - 80 - 10.8.1.1 Mozilla/4.0+(compatible;+MSIE+6.0;+Wi
ndows+NT+5.1;+SV1;+.NET+CLR+1.1.4322) 404 0 3
2006-05-31 19:59:06 W3SVC508294276 10.8.1.39 GET /index.html/neverthere.pl - 80 - 10.8.1.1 Mozilla/4.0+(compatible;+MSIE+6
.0;+Windows+NT+5.1;+SV1;+.NET+CLR+1.1.4322) 404 0 3
2006-05-31 19:59:06 W3SVC508294276 10.8.1.39 GET /index.html/cmd.pl - 80 - 10.8.1.1 Mozilla/4.0+(compatible;+MSIE+6.0;+Windo
ws+NT+5.1;+SV1;+.NET+CLR+1.1.4322) 404 0 3
```

Looking at the first few lines of the log file, you can determine the program and version that created the file and the start and end dates of the file. This can be helpful if we know the approximate time that the attack occurred.

If we know the attack happened on a Monday and the log is from the previous Wednesday, it may have little if any evidence value.

IP Search

Since this log is from the server that was attacked, searching for the server's IP address would not be useful since each entry should have that IP address in it.

If you know the IP address of the attacker, search for that, but in a NAT environment the IP could have been used or reused by another user in this same log.

It would be better to save IP searches to the end so that you can search for specific IP and times together.

Example Log, continued

String Search

We know that the attacker used an administrative account to log into the server. In order to do this, they would have entered a username and password. Let's start there.

Searching for Password

Thinking about our keywords, we will start with the possibility that Password is a good clue. We can perform a search by typing

```
grep 'pass' IIS5211_6.txt
```

This will search for the string of characters pass in every line of the log file IIS5211_6.txt. This query should return approximately 73 lines of found text.

At the top of the return you should see something like the image below. You may have to widen your command window to get the lines to look the same.

```
2006-05-31 19:59:19 W3SVC508294276 10.8.1.39 GET /index.html/password.pl - 80 - 1
2006-05-31 19:59:21 W3SVC508294276 10.8.1.39 GET /index.html/passwords.pl - 80 -
2006-05-31 19:59:45 W3SVC508294276 10.8.1.39 GET /index.html/password.php - 80 -
2006-05-31 19:59:46 W3SVC508294276 10.8.1.39 GET /index.html/passwords.php - 80 -
2006-05-31 20:00:10 W3SVC508294276 10.8.1.39 GET /index.html/password.sh - 80 - 1
2006-05-31 20:00:10 W3SVC508294276 10.8.1.39 GET /index.html/passwords.sh - 80 -
2006-05-31 20:00:34 W3SVC508294276 10.8.1.39 GET /index.html/password.py - 80 - 1
2006-05-31 20:00:36 W3SVC508294276 10.8.1.39 GET /index.html/passwords.py - 80 -
2006-05-31 20:00:59 W3SVC508294276 10.8.1.39 GET /index.html/password.tgz - 80 -
2006-05-31 20:01:01 W3SVC508294276 10.8.1.39 GET /index.html/passwords.tgz - 80 -
2006-05-31 20:01:24 W3SVC508294276 10.8.1.39 GET /index.html/password.tar - 80 -
2006-05-31 20:01:26 W3SVC508294276 10.8.1.39 GET /index.html/passwords.tar - 80 -
2006-05-31 20:01:49 W3SVC508294276 10.8.1.39 GET /index.html/password.tar.gz - 80
2006-05-31 20:01:51 W3SVC508294276 10.8.1.39 GET /index.html/passwords.tar.gz - 8
2006-05-31 20:02:15 W3SVC508294276 10.8.1.39 GET /index.html/password.asp - 80 -
2006-05-31 20:02:15 W3SVC508294276 10.8.1.39 GET /index.html/passwords.asp - 80 -
2006-05-31 20:02:40 W3SVC508294276 10.8.1.39 GET /index.html/password.aspx - 80 -
2006-05-31 20:02:40 W3SVC508294276 10.8.1.39 GET /index.html/passwords.aspx - 80
2006-05-31 20:03:05 W3SVC508294276 10.8.1.39 GET /index.html/password.doc - 80 -
2006-05-31 20:03:06 W3SVC508294276 10.8.1.39 GET /index.html/passwords.doc - 80 -
2006-05-31 20:03:30 W3SVC508294276 10.8.1.39 GET /index.html/password.exe - 80 -
2006-05-31 20:03:30 W3SVC508294276 10.8.1.39 GET /index.html/passwords.exe - 80 -
2006-05-31 20:03:55 W3SVC508294276 10.8.1.39 GET /index.html/password.cmd - 80 -
2006-05-31 20:03:56 W3SVC508294276 10.8.1.39 GET /index.html/passwords.cmd - 80 -
```


Example Log, continued

Searching for Password, continued

What we notice when looking at this view is a series of attempts to guess the password file for the system. This is not normal network traffic and is the first clue of one of the methods attempted by an attacker.

Farther down the list we find an attempt on the login page in the admin directory.

```
2006-05-31 20:46:09 W3SVC508294276 10.8.1.39 GET /admin/login.asp username=asdasd&password=
2006-05-31 20:47:39 W3SVC508294276 10.8.1.39 GET /admin/login.asp username=test&password=t
2006-05-31 20:49:52 W3SVC508294276 10.8.1.39 GET /admin/login.asp username=testFUZZCTRL&pc
2006-05-31 20:49:54 W3SVC508294276 10.8.1.39 GET /admin/login.asp username=testTo%20decryp
```

The attacker is trying different account names and password combinations.

Suspicious Text

As you look through some of the next lines that were returned, you may find a line with this information in the line:

```
2006-05-31 20:49:52 W3SVC508294276 10.8.1.39 GET
/admin/login.asp username=testFUZZCTRL&password=testpass
80
```

The text FUZZCTRL should be suspicious to you. It might be a legitimate username or password or it might not. If you search the Internet for FUZZCTRL you can find reference to a vulnerability scanning Web proxy called Suru. The manual for Suru is a free PDF download and, if viewed, you will see examples of a server attack that match the lines shown in this log.

Response

You can now correlate a known tool with an IP address and a time frame, allowing you to now proceed with a plan to contact the ISP for the domain that the attacker is coming from. You can then follow your agencies policy for contacting, serving preservation letters and obtaining warrants for information on the attacker if needed.

Lesson 3 – Automated Log Analysis Tools

Introduction

There are not many automated tools that allow you to search log files. Most require complex programming and setup prior to use with every case. We will now look at Sawmill which is one of the better tools on the market.

Purpose of this Lesson

The purpose of this lesson is to introduce the automated log analysis tool Sawmill.

Objectives

After completing this lesson, you will be able to:

- Install and configure the Sawmill program
- Describe the function and use of the Sawmill program

In this Lesson

The following table shows the contents of this lesson.

Topic	See Page
What is Sawmill?	14-30
Installing Sawmill	14-31
Network Log Analysis Using Sawmill	14-38

What is Sawmill?

Introduction

When analyzing network text logs for signs of an intrusion, a network intrusion analyst must quickly parse those logs to locate the data that correlates to the intrusion. Sawmill is a tool that will assist the analyst in parsing network text logs and organizing the logs into an easy-to-read report.

Capabilities of Sawmill

Sawmill can process various text logs generated by a variety of network security devices. Sawmill also converts the text log to a cross-linked report that allows an analyst to customize the report according to the output requirements.

Download Information

Sawmill can be purchased and downloaded from the following website:

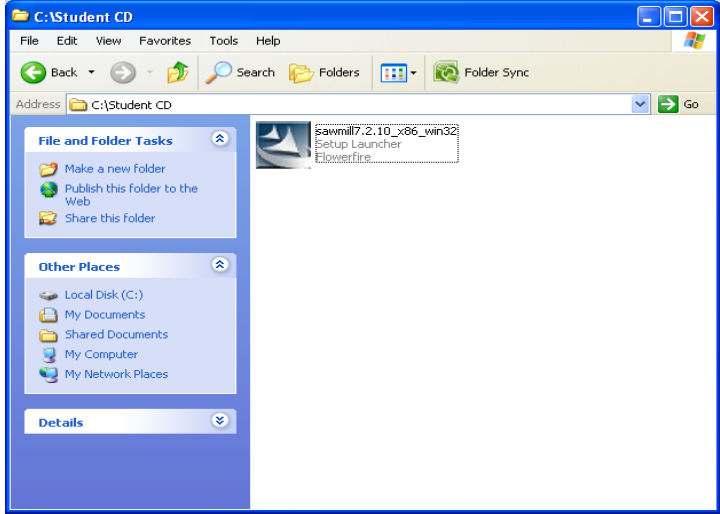
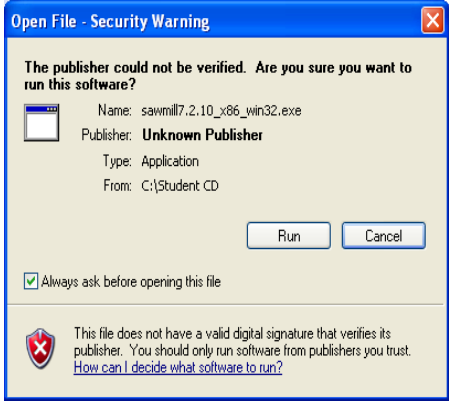
<http://www.sawmill.net>

The initial download and installation comes with a 30 day, unlimited profile license.

Installing Sawmill

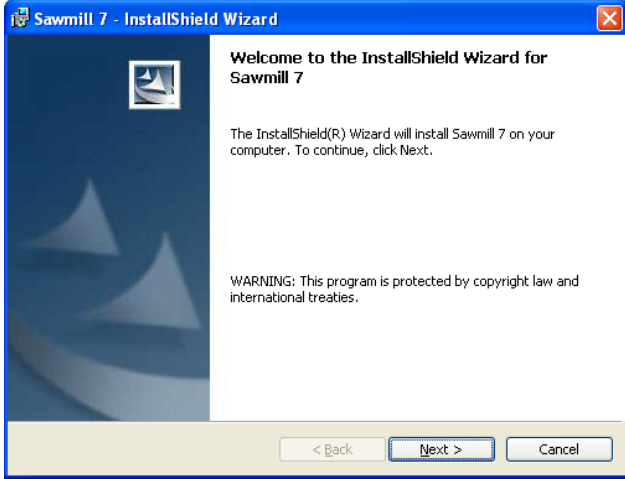

Procedure: Installing and Configuring Sawmill

Use the following procedure to install and configure Sawmill in a Windows environment.

Step	Action
1	From the system desktop, double-click on the My Computer icon.
2	<p>In the My Computer window, navigate to the location of the sawmill setup executable and double click on the Sawmill (7.2.11_x86_win32) icon.</p> 
3	<p>If you receive the Unknown Publisher warning, click on Run.</p> 


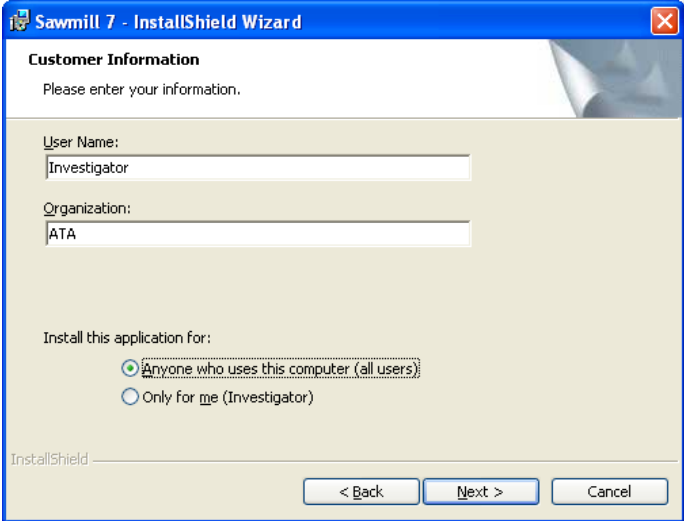
Installing and Configuring Sawmill, continued

Procedure: Installing and Configuring Sawmill, continued

Step	Action
4	<p>At the InstallShield Wizard, click Next to install Sawmill.</p> 
5	<p>Read the End User License Agreement and accept the license. Click Next.</p> 

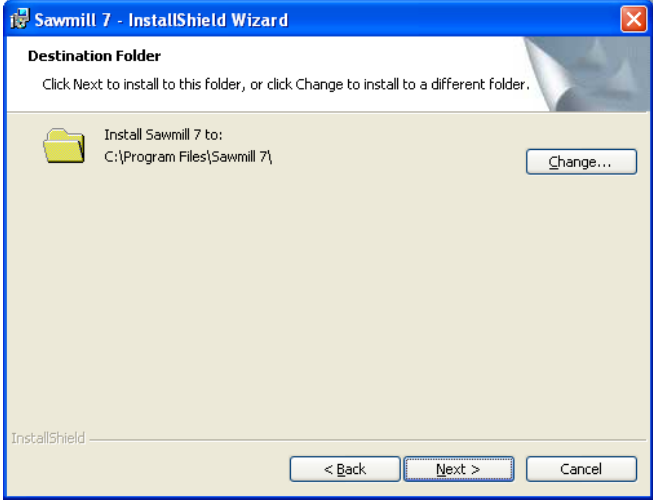
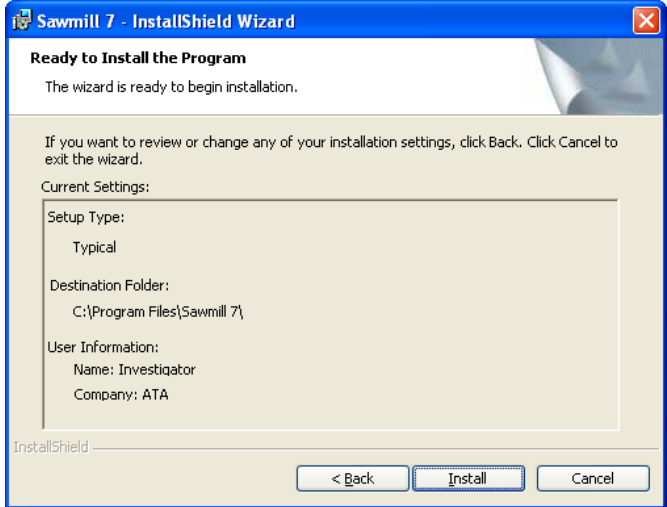
Installing and Configuring Sawmill, continued

Procedure: Installing and Configuring Sawmill, continued

Step	Action
6	<p>At the Release Notes window, click Next.</p> 
7	<p>At the Customer Information window, enter your personal details, ensure “Anyone who uses this computer (all users)” is selected, and then click Next.</p> 

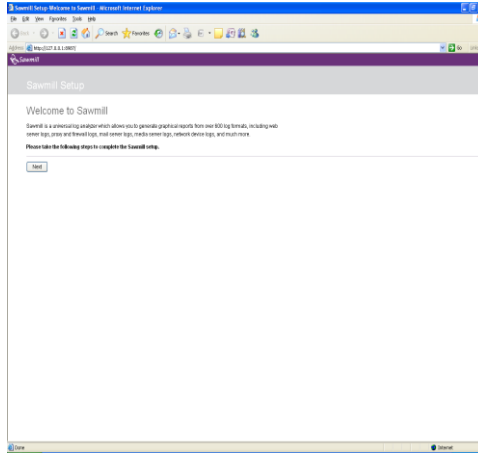
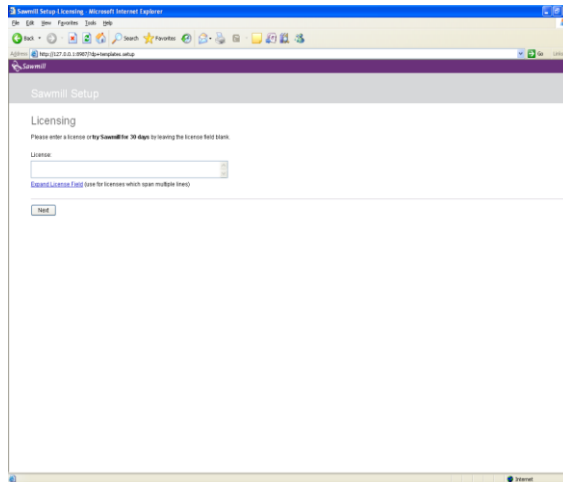
Installing Sawmill, continued

Procedure: Installing and Configuring Sawmill, continued

Step	Action
8	<p>At the Destination Folder window, click Next to install Sawmill in the C:\Program Files\Sawmill 7 folder.</p> 
9	<p>At the Ready to Install the Program window, click Install to install Sawmill.</p> 

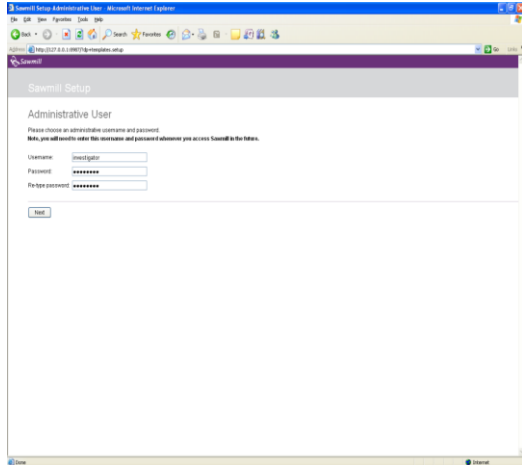
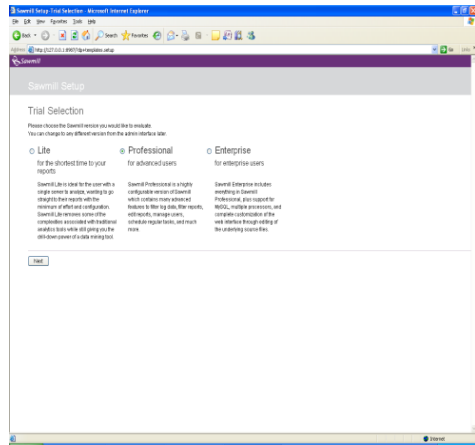
Installing Sawmill, continued

Procedure: Installing and Configuring Sawmill, continued

Step	Action
10	<p>The Sawmill setup Web interface will appear. Click Next to proceed.</p> 
11	<p>On the Sawmill Licensing setup screen, click Next. Use the 30 day trial license or enter your purchased license number.</p> 

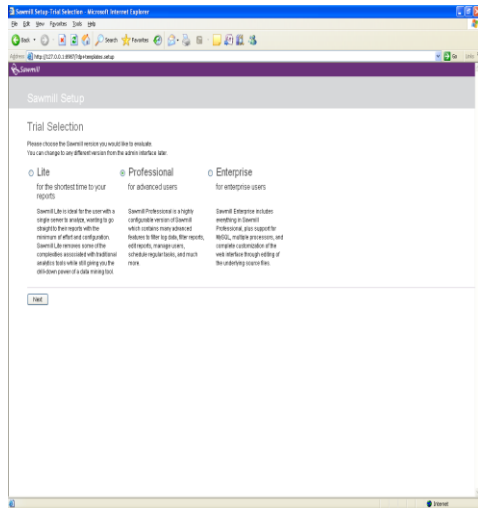
Installing Sawmill continued

Procedure: Installing and Configuring Sawmill, continued

Step	Action
12	<p>On the Administrative User setup screen, type the Username and Password and then click Next.</p> 
13	<p>On the Trial Selection Screen, select Professional and then click Next.</p> 

Installing Sawmill, continued

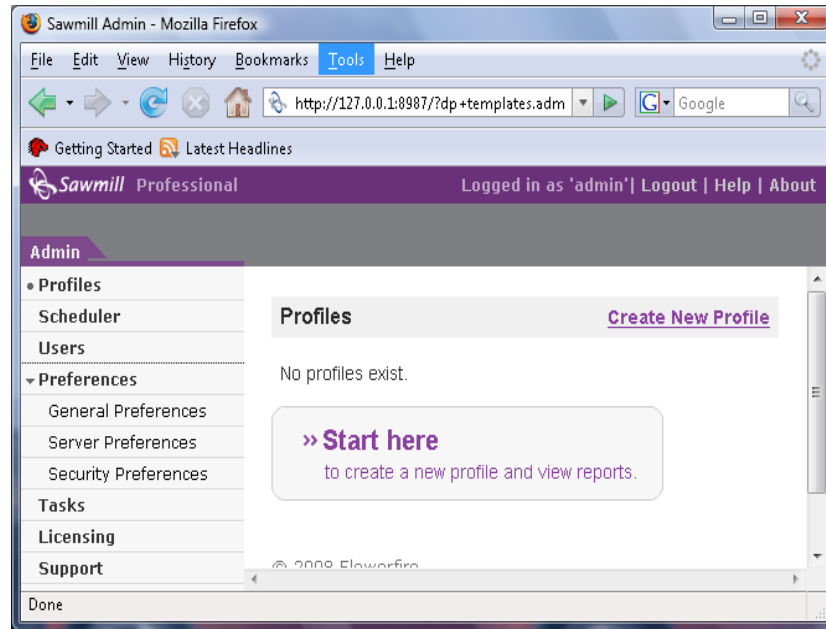
Procedure: Installing and Configuring Sawmill, continued

Step	Action
14	<p>On the Automated Feedback Agent screen, un-check the checkbox and then click Next.</p> 
15	<p>On the Complete Setup screen, click Finish to complete the setup.</p>

Network Log Analysis using Sawmill

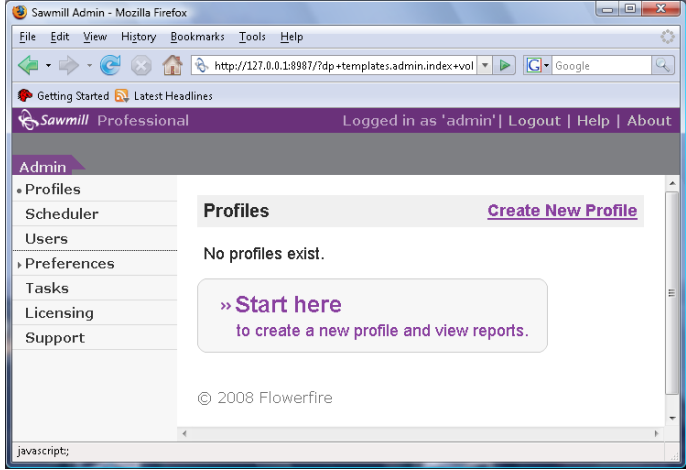
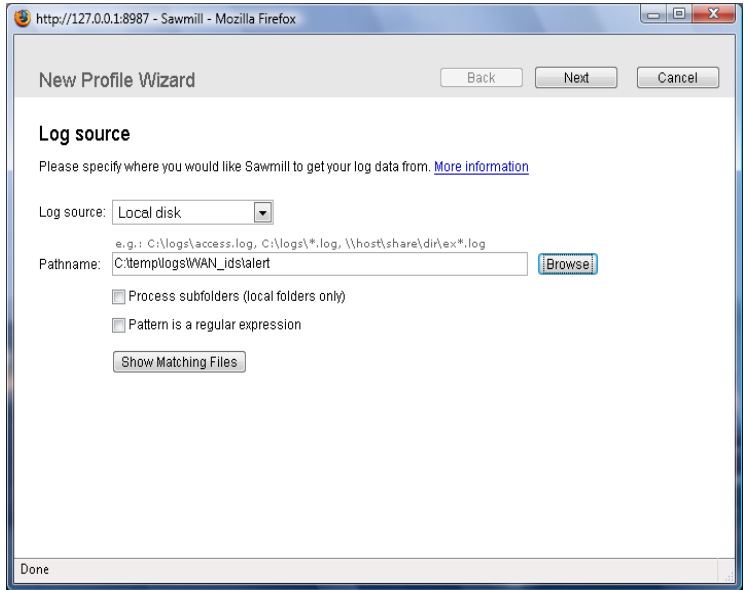
The Administrative Interface

Below is a screenshot of Sawmill's Administrative interface.



Network Log Analysis using Sawmill, continued

Procedure: In this procedure, you create a report profile and use it to parse and sort the selected text log for requested information:
Creating a Report Profile

Step	Action
1	<p>From the Administrative screen, select “Create New Profile.”</p> 
2	<p>In the New Profile Wizard, select Log Source: Local Disk. For Pathname, click Browse. Navigate to the log file location and select the log file. Click Next.</p> 

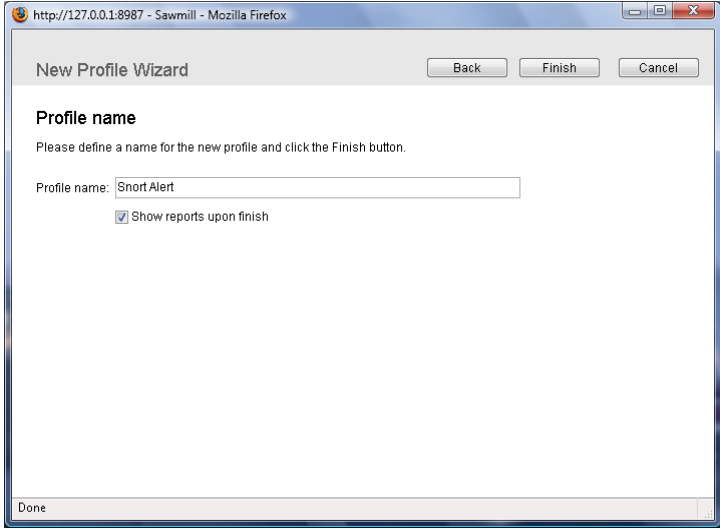
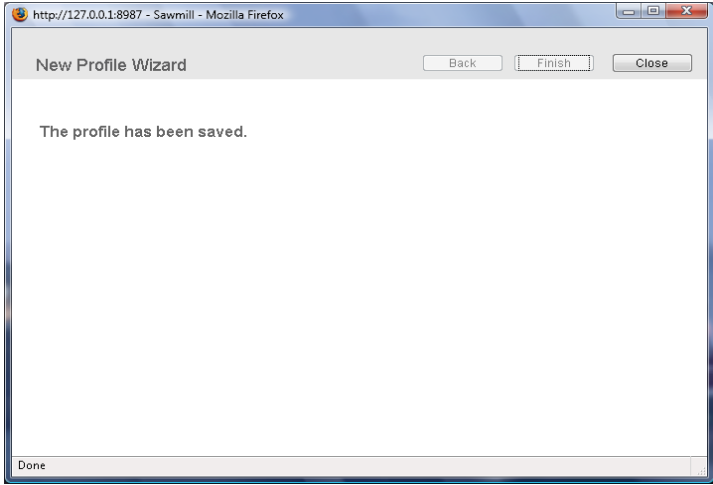
Network Log Analysis using Sawmill, continued

Procedure: Creating a Report Profile continued

Step	Action
3	<p>Sawmill automatically attempts to detect the log format and display the results for user selection. Select the log format. At the bottom of the display ensure that “Continue with the above detected log formats” is selected. Click Next.</p> More information'. A list box contains four options: 'MM/DD-HH:MM:SS Timestamp', 'Snort 2 Log Format (syslog required)', 'Snort Log Format (standalone, mm/dd dates)' (which is selected), and 'Snort Log Format (syslog required)'. Below the list box are two radio buttons: 'Continue with one of the above detected log formats (recommended)' (which is selected) and 'Choose a different log format on the next wizard page.' followed by explanatory text and a 'click here' link. A 'Done' button is at the bottom left." data-bbox="413 325 843 616"/>

Network Log Analysis using Sawmill, continued

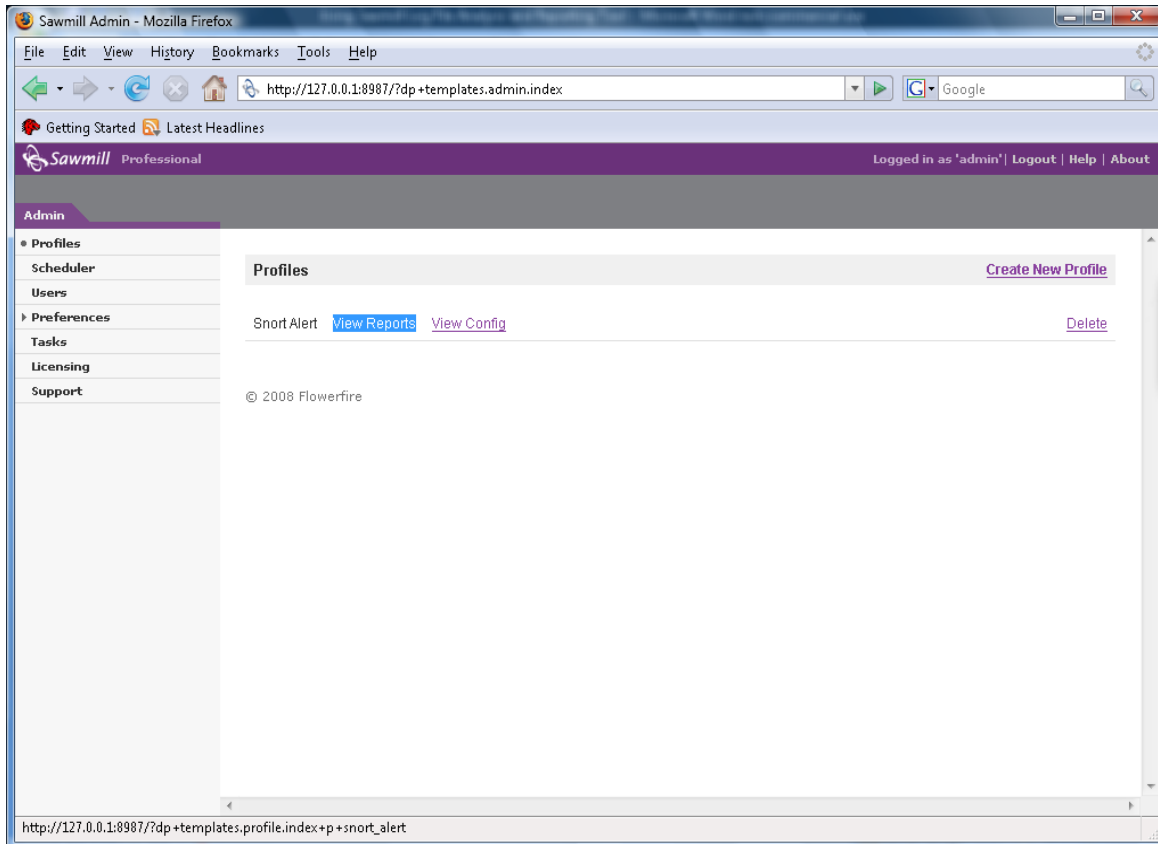
Procedure: Creating a Report Profile continued

Step	Action
4	<p>In the “Profile name” type the desired name of the profile. Ensure the “Show reports upon finish” checkbox is selected. Click Finish.</p> 
5	<p>The profile is saved in the final screen of the New Profile wizard. Select Close.</p> 

Network Log Analysis using Sawmill, continued

The Report Environment

Once a report profile has been generated, you can load the selected report by selecting “View Reports” from the Administrative screen.

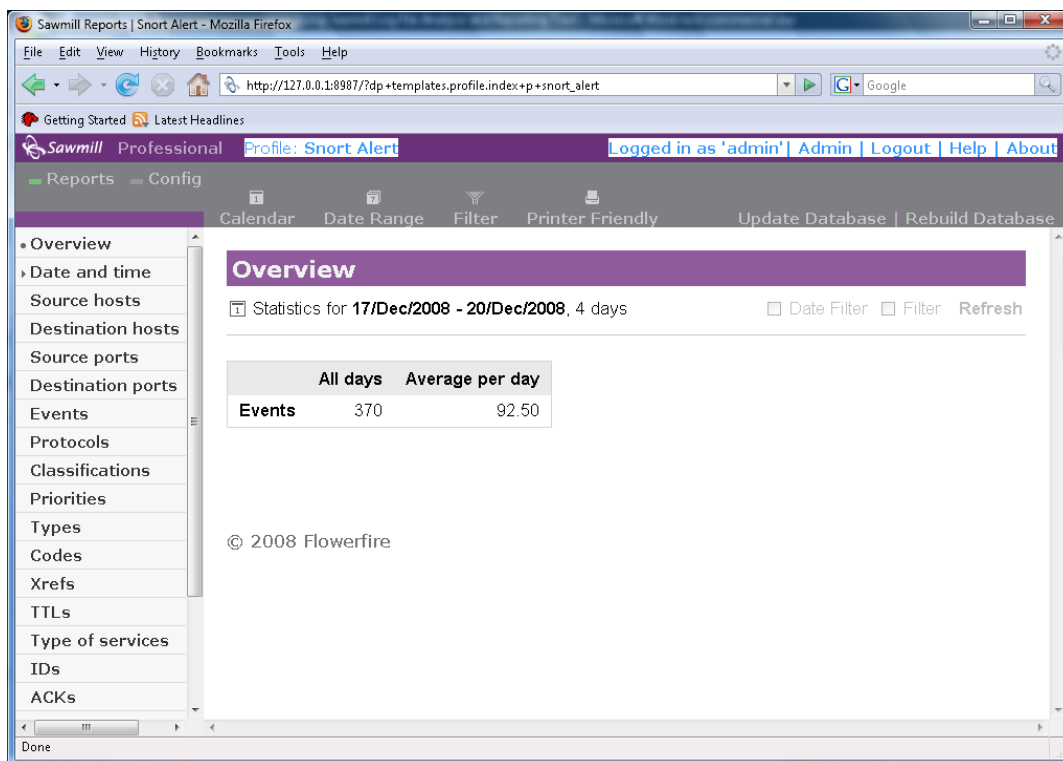


Network Log Analysis using Sawmill, continued

Report Header

The header of the report contains the following information:

- Profile name – The name of the active profile which is being displayed
- Admin link – Link to the administrative functions, such as profile lists
- Logout – A link to log out of Sawmill
- Help – Help documentation

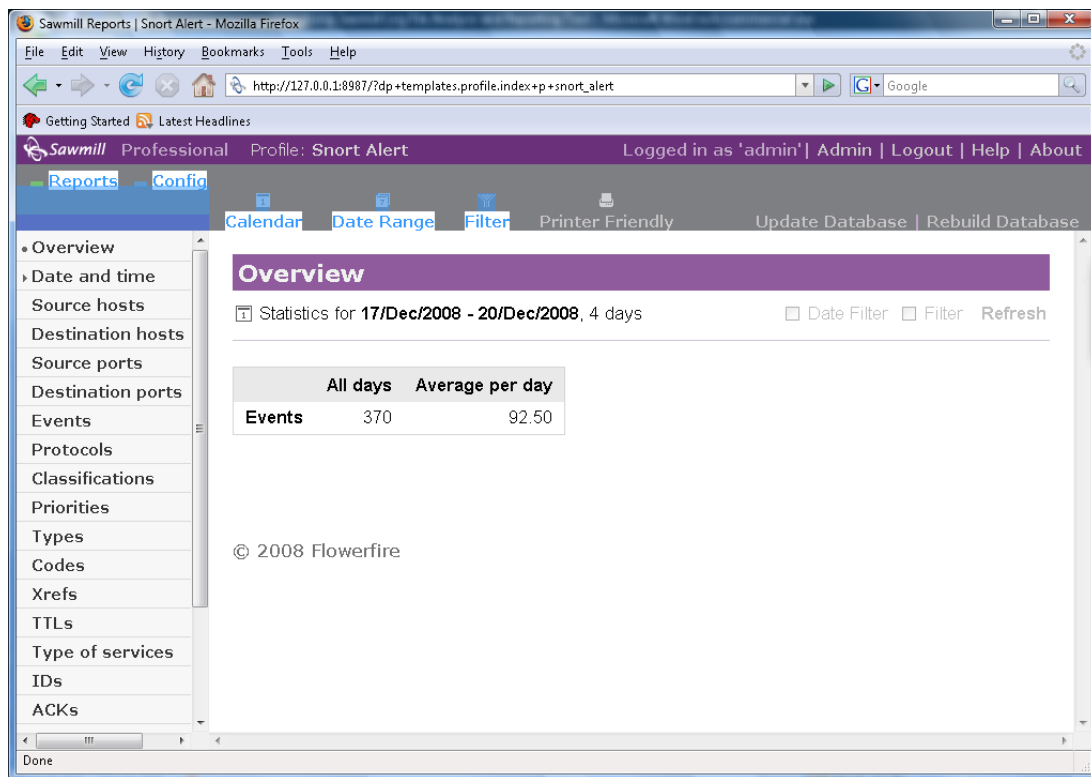


Network Log Analysis using Sawmill, continued

Report Toolbar

Below the report's header is the report's toolbar. This toolbar contains the following links:

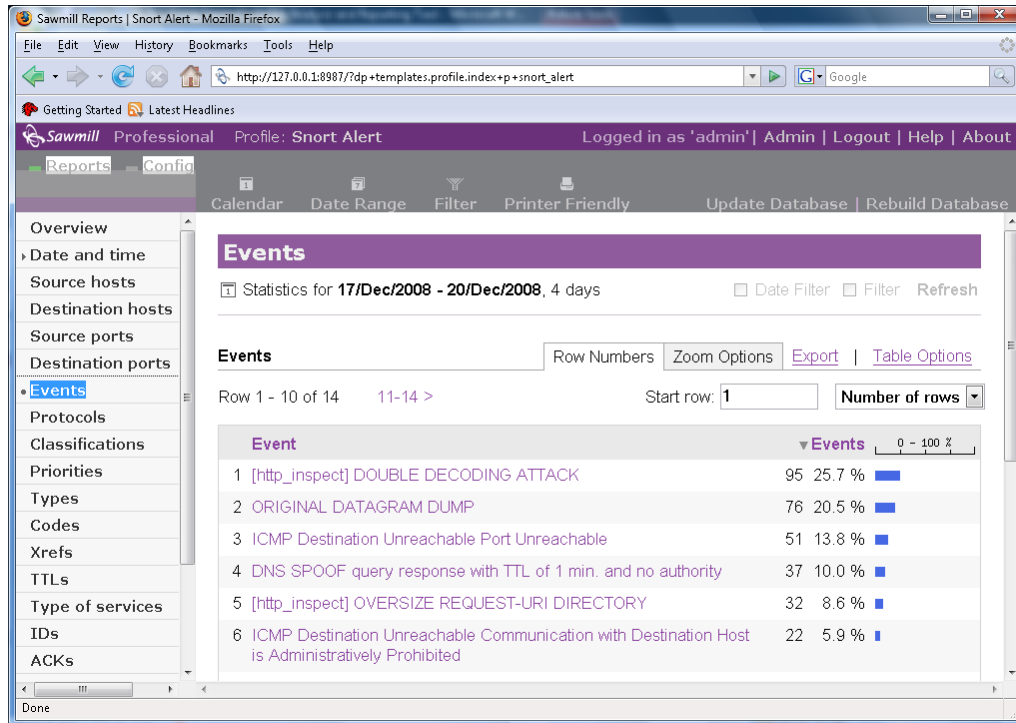
- Reports - Used to access other loaded reports from the current reports view
- Config – Allows you to change profile options
- Calendar – Date/time filter can be set to view a single day, month or year
- Date Range – A range of days can be selected to use as the date/time filter
- Filter – Used to configure global filter options for any of the report fields. These filters dynamically affect all reports.



Network Log Analysis using Sawmill, continued

Report Menu

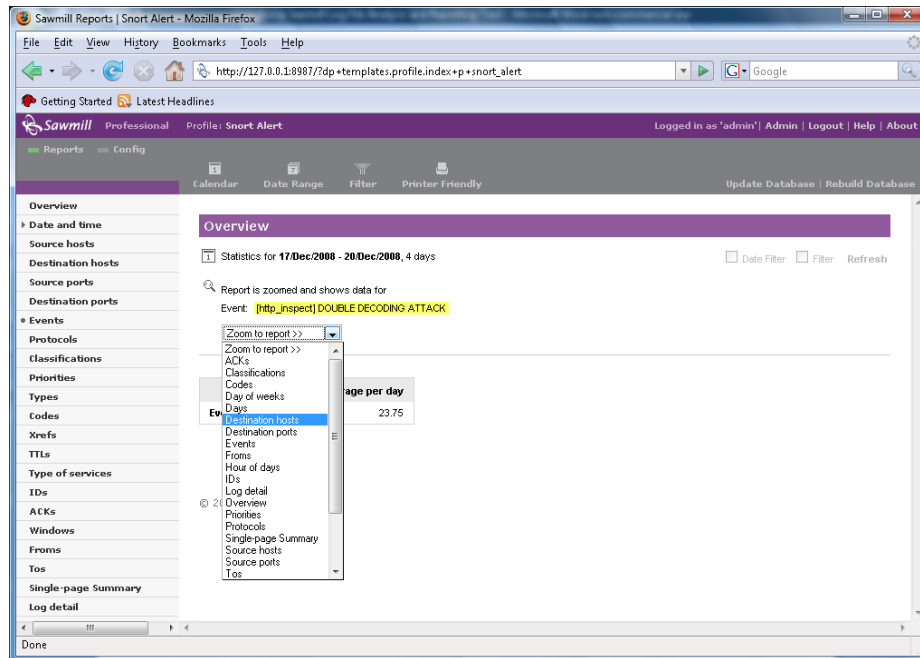
At the left of the selected report is the report's menu which lets you select different attributes of the report to view. Clicking on a report category expands or collapses the category and allows the report to zoom in on and display the selected category.



Note: In the above example, the **Events** attribute was selected to display all events.

Network Log Analysis using Sawmill, continued

Zoom To Filters Once a report attribute is selected from the report's menu, the report display can be filtered using the Zoom to Report feature.



Note: In the above example, the “Destination host” zoom filter is selected. Continuously selecting Zoom to Report filters would narrow the intended search.

Network Log Analysis using Sawmill, continued

Final Output Report (Log Detail)

Once all filters have been applied to the report, a final log detail can be generated to display all of the set attributes. To do this, click on “Log Detail” from the Zoom To Report filter box.

Sawmill Reports | Snort Alert - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://127.0.0.1:8987/?dp+templates.profile.index+p+snort_alert

Getting Started Latest Headlines

Sawmill Professional Profile: Snort Alert Logged in as 'admin' | Admin | Logout | Help | About

Reports Config

Calendar Date Range Filter Printer Friendly Update Database Rebuild Database

Overview Statistics for 17/Dec/2008 - 20/Dec/2008, 4 days

Report is zoomed and shows data for

Event: [http_inspect] DOUBLE DECODING ATTACK

Zoom to report >>

Zoom to report >>

ACKs

Classifications

Codes

Day of weeks

Days

Destination hosts

Destination ports

Events

Froms

Hour of days

IDs

Log detail

Overview

Priorities

Protocols

Single-page Summary

Source hosts

Source ports

Tos

Row Numbers Zoom Options Export Table Options

Start row: 1 Number of rows

	Destination host	Source port	Destination port	Event	Protocol	From	To
1	64.156.13.20	1716	80	[http_inspect] DOUBLE DECODING ATTACK	TCP	10.3.20.80:1716	64.156.13.20:80
2	64.236.29.63	1212	80	[http_inspect] DOUBLE DECODING ATTACK	TCP	10.3.20.80:1212	64.236.29.63:80
3	17/Dec/2008 15:07:30	10.3.20.80	64.236.29.63	1521	80	[http_inspect] DOUBLE DECODING ATTACK	TCP 10.3.20.80:1521 64.236.29.63:80
4	17/Dec/2008 16:04:34	10.3.20.80	128.241.21.163	1209	80	[http_inspect] DOUBLE DECODING ATTACK	TCP 10.3.20.80:1209 128.241.21.163:80

Done

Network Log Analysis using Sawmill, continued

Final Output Report, continued

Here is the screen of the final sorted output.

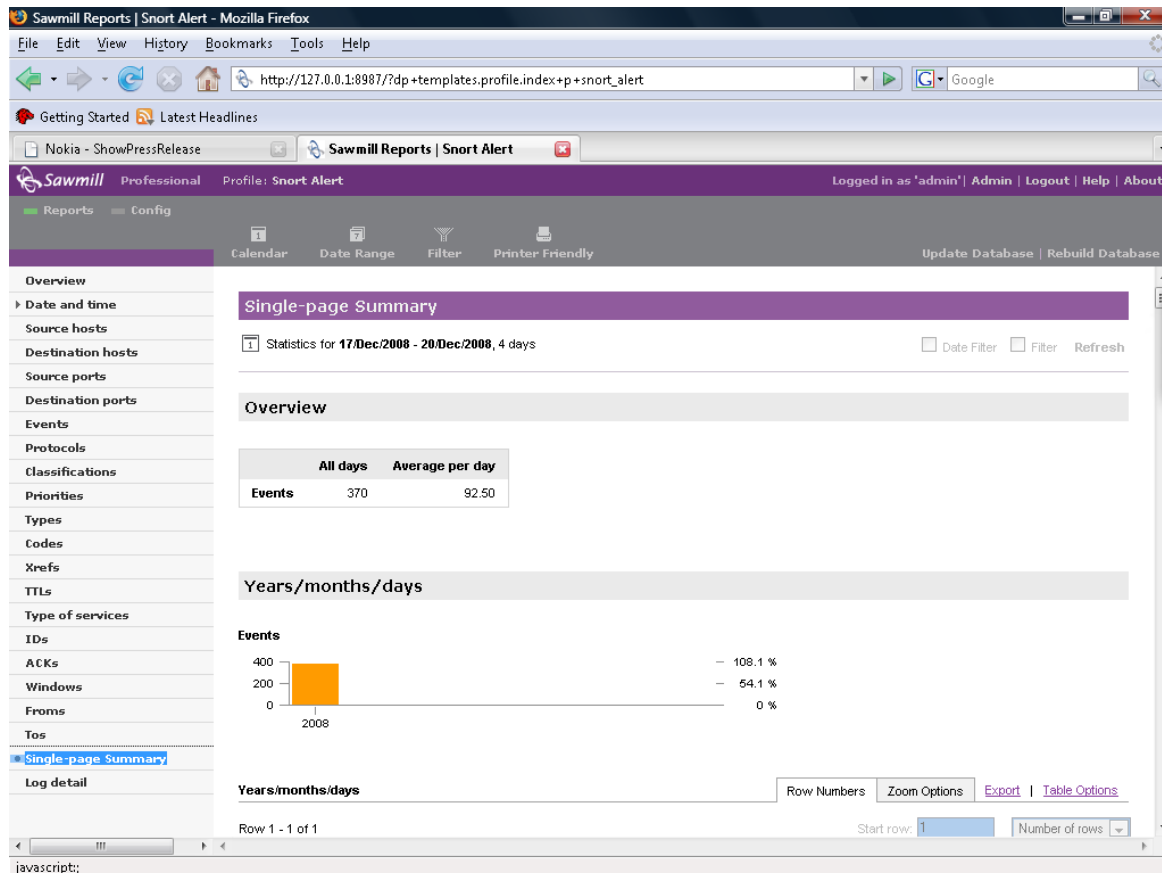
The screenshot shows the Sawmill Reports interface in a Mozilla Firefox browser window. The URL is `http://127.0.0.1:8987/?dp+templates:profile.index+p+snot_alert`. The interface includes a navigation menu on the left with options like Overview, Date and time, Source hosts, Destination hosts, Source ports, Destination ports, Events, Protocols, Classifications, Priorities, Types, Codes, Xrefs, TTLs, Type of services, IDs, ACKs, Windows, Froms, Tos, Single-page Summary, and Log detail. The main content area displays statistics for the period 17/Dec/2008 - 20/Dec/2008, 4 days. A search filter is applied for the event '[http_inspect] DOUBLE DECODING ATTACK'. The 'Log detail' section shows a table of log entries.

	Date/time	Source host	Destination host	Source port	Destination port	Event	Protocol	From	To
1	17/Dec/2008 10:53:22	10.3.20.80	64.156.13.20	1716	80	[http_inspect] DOUBLE DECODING ATTACK	TCP	10.3.20.80:1716	64.156.13.20:80
2	17/Dec/2008 13:55:11	10.3.20.80	64.236.29.63	1212	80	[http_inspect] DOUBLE DECODING ATTACK	TCP	10.3.20.80:1212	64.236.29.63:80
3	17/Dec/2008 15:07:30	10.3.20.80	64.236.29.63	1521	80	[http_inspect] DOUBLE DECODING ATTACK	TCP	10.3.20.80:1521	64.236.29.63:80
4	17/Dec/2008 16:04:34	10.3.20.80	128.241.21.163	1209	80	[http_inspect] DOUBLE DECODING ATTACK	TCP	10.3.20.80:1209	128.241.21.163:80

Network Log Analysis using Sawmill, continued

Single Page Summary

You can generate a single page summary containing all log attributes using the “Single Page Summary” category located in the report’s menu. Here is an example of a single page summary.



This page intentionally left blank.

Module 12

LiveWire Investigations

Introduction

This module will introduce you to the Wetstone LiveWire Investigator tool, commonly referred to as LiveWire. The tool is used to conduct live digital investigations. Other tools used in this module include additional Wetstone tools and open source tools that can be used during live investigations.

Purpose of this Module

The purpose of this module is to show you how to setup your workstation. You will learn how to install and use tools used in live digital investigations.

Objectives

After successfully completing this module, you will be able to:

- Properly prepare for a live digital investigation
- Use live digital investigation tools

In this Module

The following table shows the contents of this module.

Topic	See Page
Lesson 1 – Data Collection	12-3
Lesson 2 – Introduction to LiveWire	12-9
Lesson 3 – LiveDiscover	12-31
Lesson 4 –Volatile Data Analysis	12-39
Lesson 5 – Evidence Collection	12-73
Lesson 6 – Malicious Code Analysis	12-93
Lesson 7 – Alternate Data Collection Tools	12-99

This page intentionally left blank.

Lesson 1 – Data Collection

Introduction

When collecting data for any investigation it is vital that the data collection is conducted correctly. This information may contain the only source of evidence in an investigation and should be collected accurately and correctly to be admissible in court.

Purpose of this Lesson

This lesson will discuss the importance of collecting data in the proper manner.

Objectives

After completing this lesson, you will be able to:

- Discuss locating physical devices in a network environment
- Explain how to collect data to forensically clean media

In this Lesson

The following table shows the contents of this lesson.

Topic	See Page
Locating Physical Devices	12-4
Attaching Storage Equipment	12-6

Locating Physical Devices

Network Architecture

When trying to locate a particular server, the logical topology map shows how things are connected logically, but not necessarily physically. A logical map details how data flows across a network, but not how it is physically wired. Network architecture indicates where devices are physically located.

Network devices that provide a possible path for the incident are considered to be “in-line” to the investigation. As these devices have carried traffic relating to the incident, they may hold crucial information and should be properly located and analyzed.

Logical Assessment

Logical assessment involves obtaining any network topologies to get a rough estimate of where sensors can be placed for an investigation. The network topologies may not exist or be severely outdated. An investigator can update the topology through interviews or by performing a physical assessment.

Physical Assessment

Physical assessment includes tracing wire and cable to physical components on the network to create a wiring diagram. A wiring diagram shows the physical connections between devices onsite and can help determine the accuracy of the logical assessment. An investigator can use several cable testing devices, like a tone generator, to verify cable locations.

In large network environments, servers and network devices are assigned some form of inventory control, such as a bar code or unique name. It may be necessary to search through rows of server racks to locate an identification tag on the server of interest.

Performing a Physical Site Examination

You need to examine the physical site to determine the physical data paths and their relationship to the overall physical environment. Understanding these relationships provides a basis for determining what is or is not physically possible on the network. A physical site examination includes the following tasks:

- Physically locate the target host
- Physically locate the device to which the target host is connected
- Physically locate devices that fall into the path of the investigation
- Verify the network documentation (if available)

Locating Physical Devices, continued

Verifying the Network Configuration

When encountering an unfamiliar network, an investigator needs a starting point to verify the network setup and the actual location of network devices. Almost every network has a connection to the Internet or some external network. This external link is typically the best starting point to begin tracing wire.

Tracing wire is a technique used to determine how devices are physically connected to each other. If a wire cannot be traced because it is tightly bound to other cables or travels into a wall, other devices like a network tone generator can be used to determine its termination location. However, the use of some of these devices could require unplugging the cable and severing any existing connections, which could alert the suspect(s) of your presence.

Physically Locating the Target Host

To physically locate the target host, you must collect all identifying information regarding the device from your review of the network documentation and interview with the system administrator. Use this information and your physical assessment of the network to locate the device.

Physically Locating the Nearest Device

Locate the hub or switch to which the target host is connected. This can be found by:

- Using the identifying information that you obtained during your review of the network documentation
- Recording the termination location for each network-capable cable connected to the machine. This could be an RJ45 or RJ11 socket on the nearest wall, a hub or switch, or some other device. If the cable terminates at a wall socket, record that socket's ID number and locate it on the patch panel that aggregates cables for that area of the facility.

Attaching Storage Equipment

Data Storage

Many investigations result in large evidence files that must be collected to the investigator's computer. Sufficient data storage to copy and preserve evidence files is imperative to a successful investigation.

Whenever data is retrieved from the suspect machine, the data should always be redirected to a forensically clean evidence collection drive. The investigator must be sure to never save the output of an investigation to the local system's hard drive, as it may compromise the evidence as well as potentially fill the computer's disk space.

The storage equipment can connect to the collection machine by many different connection types. Some of the most common types are: USB, FireWire, and eSATA. External hard drives with these configurations come in many different capacities, and the general rule is to allocate as much disk space as possible for each investigation.

Wiping and Verification

When collecting evidence from any system, the data should be stored on a forensically clean drive. In order to be forensically clean, the evidence storage drive must be thoroughly wiped. Wiping is the process of overwriting every bit on the drive using a known character or set of characters. Once the storage drive is completely overwritten the process should be verified to ensure its success. Failure to properly complete this process can result in claims of contaminated evidence, which may jeopardize the credibility of your evidence.

Why Wipe Disks?

Wiping utilities allow you to wipe an entire hard drive so that no data is left on the drive. Wiping does not simply delete the data; it overwrites every sector on your drive with a hex character. This eliminates any possibility of previous data from another case contaminating the current case.

Attaching Storage Equipment, continued

Wiping Guidelines Regardless of the tool used, always adhere to the following guidelines:

1. Clearly identify media to be wiped and segregate it from other media.
2. Have only essential media in the system during wiping operations.
3. Ensure the correct media has been selected before executing any wipe utility.
4. Remove wiped media from the machine immediately after wiping and store separately.
5. Annotate in your case notes that you wiped the media prior to its use.
6. Label media with software version and command line used (with all options).

This page intentionally left blank.

Lesson 2 – Introduction to LiveWire

Introduction

Unlike traditional dead box investigations, live digital investigations blur the line between network intrusions and evidence collection.

Purpose of this Lesson

The purpose of this lesson is to prepare you for a live investigation. You will learn the terms and concepts associated with a live investigation and prepare a system to conduct live investigations of a networked system.

Objectives

After successfully completing this lesson, you will be able to:

- Explain the basic concepts of a live digital investigation
- Install, update, and setup LiveWire Investigator
- Install and update LiveDiscover

In this Lesson

The following table shows the contents of this lesson.

Topic	See Page
Live Digital Investigations	12-10
LiveWire Installation	12-13
LiveDiscover Installation	12-14
Updating LiveWire	12-16
Updating LiveDiscover	12-17
LiveWire Initial Setup	12-19

Live Digital Investigations

Definition

When performing traditional computer seizures and then examining the media from a system that has had power removed, you are investigating only the data at rest on the media prior to the power being removed. This is also known as dead box forensics.

Live digital investigations are performed on running systems prior to the removal of power.

Why Live?

An increasing amount of memory resident programs and utilities revert to an obfuscated or encrypted state at power off. Therefore, it is sometimes necessary to seize information from the volatile areas of a computer *before* the plug is pulled.

If you are dealing with a time-sensitive case, you may also want to perform a live digital investigation in order to gather information and evidence as quickly as possible.

If you can connect to the system via a network and have sufficient access to the system in question, it is possible that you can capture the memory contents, process list, and other volatile resources before pulling the plug

Here are some possible reasons for choosing to perform a live investigation:

- Rapid response requires remote investigation
- Network size limits flexibility
- Encrypted file system requires live capture
- System of interest is mobile
- Commercial system cannot be shutdown

Live Digital Investigations, continued

How It Works

LiveWire is a complex series of scripts, programs and tools that combine to give you insight into a network and the various machines on that network.

At a deeper level, LiveWire uses a customized version of Apache Web server on your investigation system to provide you menus, displays and reports in a graphic user interface.

Additionally, LiveWire includes an embedded version of Gargoyle malware detection software, also a product of Wetstone.

Note: The use of LiveWire requires an account on a target system with administrative privileges to access and retrieve data.

LiveWire uses a Connect-Act-Disconnect model for communicating with a host on the network. For example, if you want to view all running processes on the system, LiveWire will use the administrative account and password you supplied by a local administrator when starting the investigation. The software will log into the system, obtain a list of the running processes and log out of the system.

Risks

If you are trying to access a system covertly, you should be aware that some tasks performed by LiveWire can be resource intensive and thereby noticed if a user is on the system at the time.

In some cases, the user may think that the network is slow and ignore the change in machine behavior, or if the user is fairly computer literate, he or she may be able to detect the activity by monitoring the Task Manager. Some users react negatively to such activity and may pull the network connector and start hiding evidence.

It is often said that if you are investigating a live machine clandestinely you should wait until the suspect is away from the system before running most queries in order for your activities to remain unnoticed.

Live Digital Investigations, continued

Workstation Setup You will want a reasonable quality system for performing live investigations. As with any forensic examination, ensure that any media onto which you save evidence is forensically wiped prior to use.

The minimum system requirements for LiveWire are:

- Microsoft Windows XP
- 100 MB of free disk space
- 128 MB RAM
- Pentium 300 Mhz
- Network Interface Card
- CD-ROM Drive for installation
- VGA Resolution Monitor
- Mouse

Choosing Hardware

For a computer system used for investigative purposes, bigger and faster is almost always the best choice. You should try to determine the scope of the investigation and allocate storage and processor space that will meet or exceed the expected results.

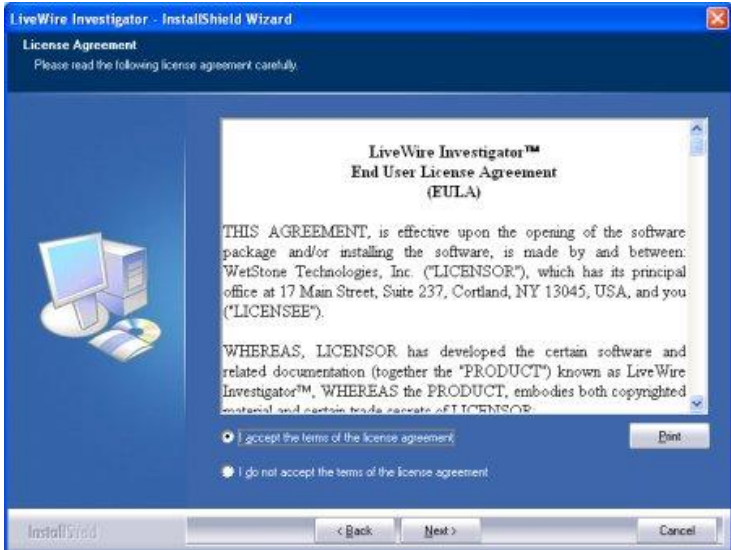
LiveWire Installation

Installation Overview

Before an application can be used on any system, it must first be successfully installed.

Procedure: LiveWire Installation

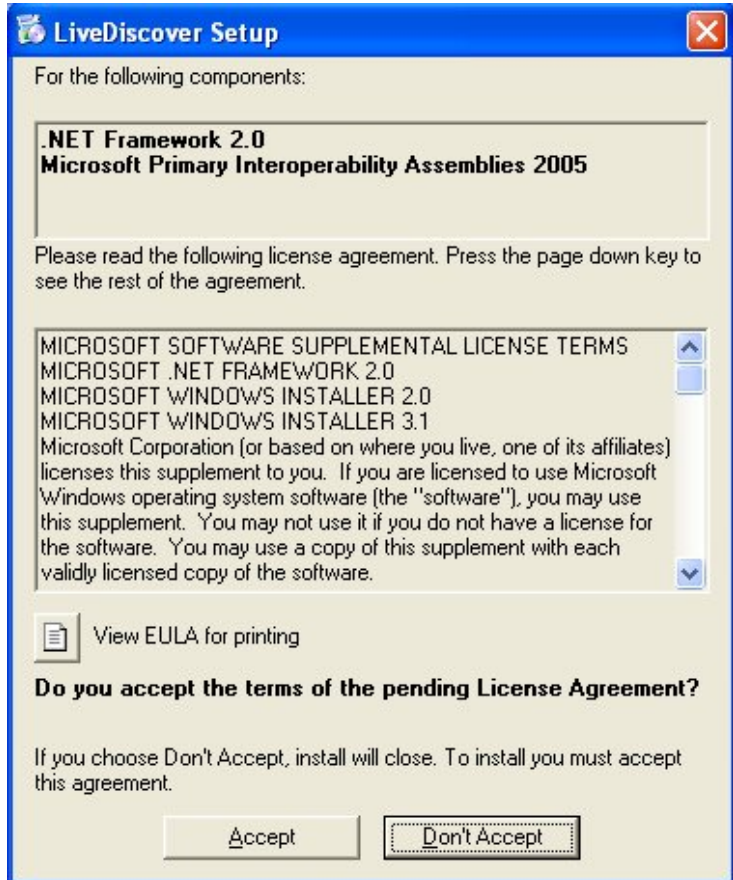
Use the following steps to install LiveWire onto a system running the Windows XP operating system.

Step	Action
1	Insert the LiveWire Investigator CD into the CD/DVD tray. Browse to its designated drive letter to view the contents of the CD.
2	Double click the livewireinstaller.exe file to execute the file, and click Next .
3	Read the EULA carefully and select "I accept the terms of the license agreement." Click Next to continue.
	
4	Click Install to begin installation.
5	Click Finish to complete the installation process.
6	Remove the LiveWire Investigator Install CD from the CD/DVD drive tray.

LiveDiscover Installation


Procedure: Installation of LiveDiscover

Use the following steps to install LiveDiscover onto a system running the Windows XP operating system.

Step	Action
1	Insert the LiveDiscover CD into the CD/DVD tray. Browse to its designated drive letter to view the contents of the CD.
2	<p>Double click the setup.exe file to execute the file. A EULA for the .NET Framework 2.0 and Microsoft Primary Interoperability Assemblies 2005 will be displayed. Carefully read the agreement, and then click Accept to continue.</p> 

LiveDiscover Installation, continued

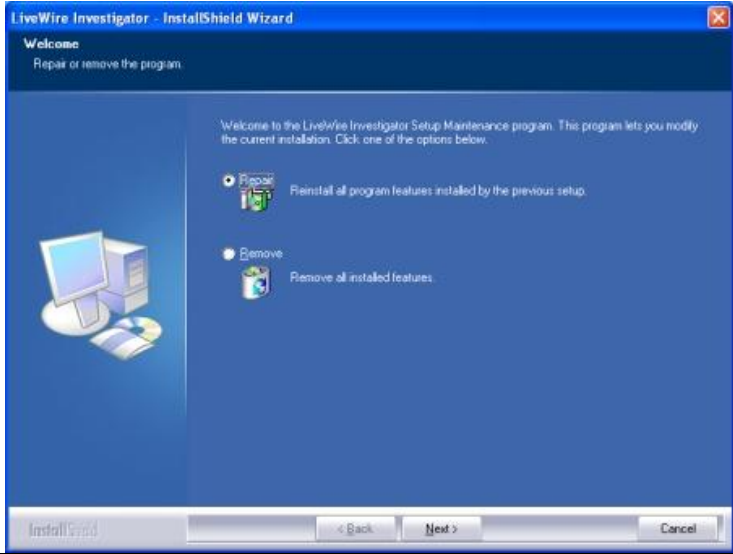
Procedure: Installation of LiveDiscover, continued

Step	Action
3	Next, the Microsoft .NET Framework 2.0 will be installed. The process will continue after the installation is complete.
4	<p>After the .NET Framework is completed, the LiveDiscover setup wizard will be displayed. Click Next.</p> 
5	Read the license agreement carefully. Select “I Agree” and click Next to continue.
6	Next, you will have the opportunity to select an installation directory. Click Next to accept the default settings.
7	Confirm Installation Page. Click Next to begin installation.
8	Once installation has completed, click Close to exit installer.
9	Remove the LiveDiscover media from the CD/DVD drive tray.

Updating LiveWire

Procedure:
Updating
LiveWire

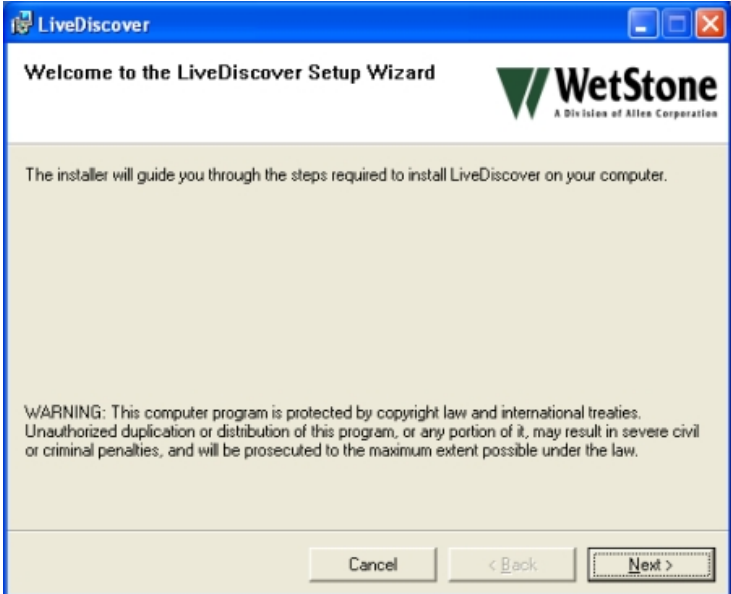
Use the following steps to update the LiveWire application.

Step	Action
1	Obtain the LiveWire update files on CD or other source. Click the Livewireintaller.exe to begin the update.
2	<p>You will be prompted to repair or remove. Choose Repair, and click Next to begin the update.</p> 
3	Once the update has completed, click Finish to exit installation wizard.

Updating LiveDiscover

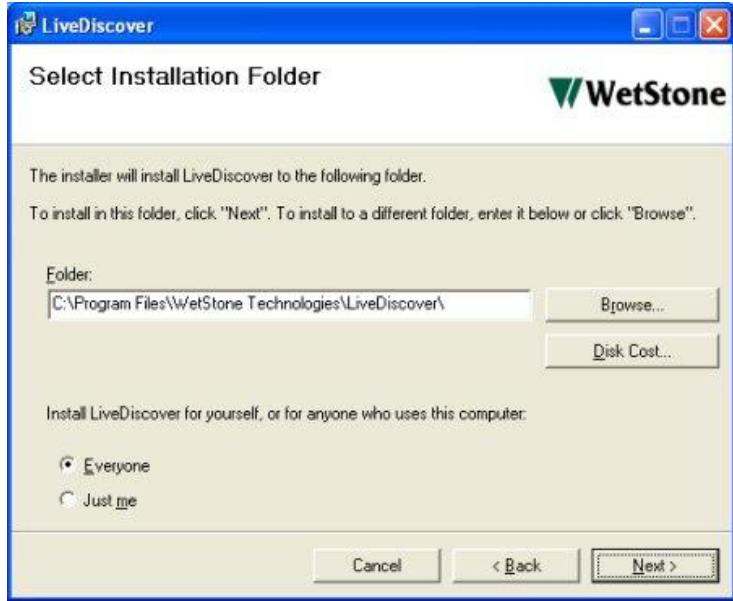
Procedure:
Updating
LiveDiscover

Use the following steps to update the LiveDiscover application.

Step	Action
1	Obtain the LiveWire update files on CD or other source. Click the Setup.exe .
2	When the Welcome screen is displayed, click Next to continue. 

Updating LiveDiscover, continued

Procedure: Updating LiveDiscover, continued

Step	Action
3	Review the licensing agreement, and select "I Agree." Click Next to continue.
4	<p>Verify that the installation directory is the current location of the LiveDiscover install directory. Click Next.</p> 
5	Confirm Installation Page. Click Next to begin installation.
6	Once installation has completed, click Close to exit installer.

LiveWire Initial Setup

Introduction

Before an investigation can be performed with LiveWire Investigator, it must be properly setup. This includes setting up the Administrator account and creating an Investigator account. The Administrator account is used to manage the investigator accounts. The investigator accounts are used to conduct the actual investigation.


Procedure: LiveWire Setup

Use the following steps to prepare LiveWire Investigator.

Step	Action
1	<p>To start the LiveWire services choose: Start > All Programs > LiveWire > Launch LiveWire. A box will pop up to verify that the services are starting.</p> 


LiveWire Initial Setup, continued

Procedure: LiveWire Setup, continued

Step	Action
2	<p>Once the services are started a new browser window will be opened to the address https://localhost/. This may cause a Security Alert box to be displayed. Click Yes to proceed.</p> 

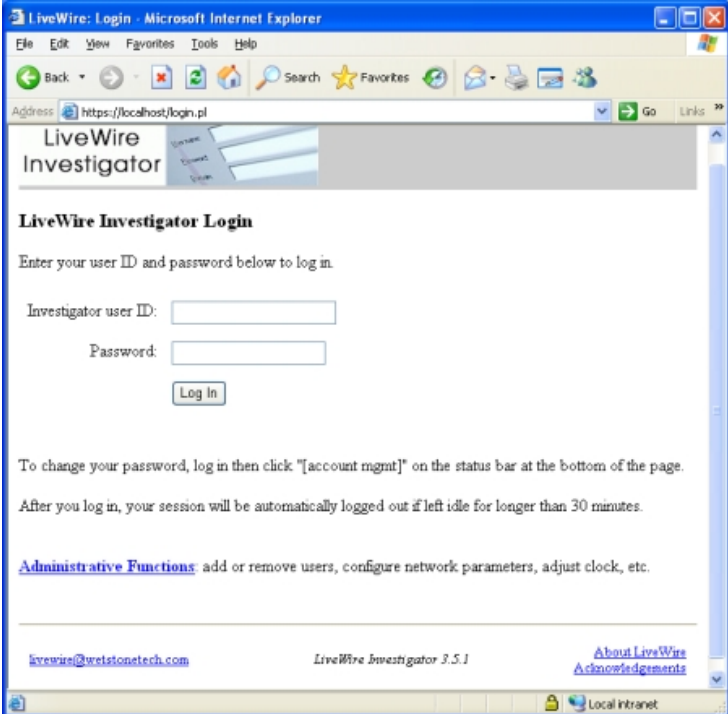
LiveWire Initial Setup, continued

Procedure: LiveWire Setup, continued

Step	Action
3	<p>The first page to open will be the license agreement page. Read the agreement and click “I Agree” to continue.</p> 

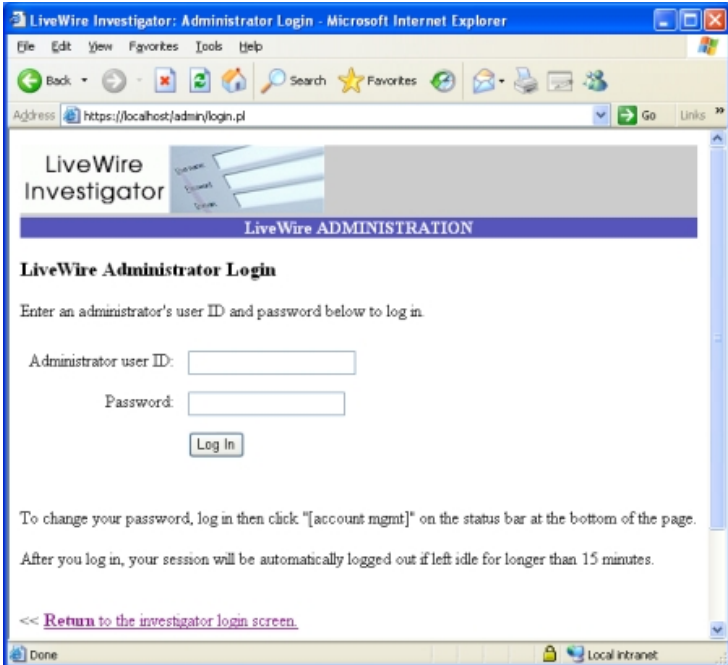
LiveWire Initial Setup, continued

Procedure: LiveWire Setup, continued

Step	Action
4	<p>The default LiveWire Investigator Login screen will be displayed. An investigator user account must be created the first time LiveWire is run. To create the account, click on Administrative Functions at the bottom of the page.</p> 

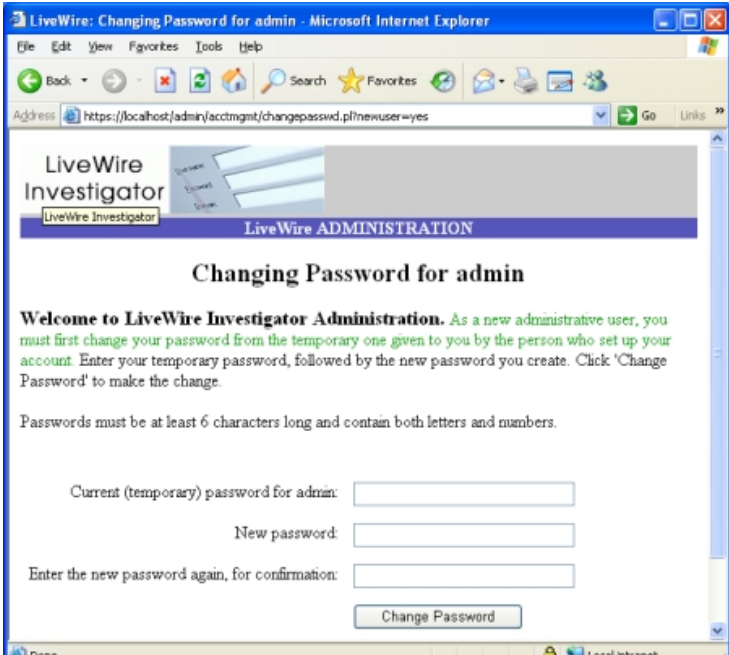
LiveWire Initial Setup, continued

Procedure: LiveWire Setup, continued

Step	Action
5	<p>The next screen will be the Administrator Login page. Login with the default setting:</p> <p>Administrative user ID: admin Password: wetstone</p> 


LiveWire Initial Setup, continued

Procedure: LiveWire Setup, continued

Step	Action
6	<p>The first time the administrator logs in, the password must be changed. The password must meet certain criteria. It must contain letters and numbers or other symbols. Input the current password of “wetstone” and then enter a new password. Then click on Change Password.</p> 

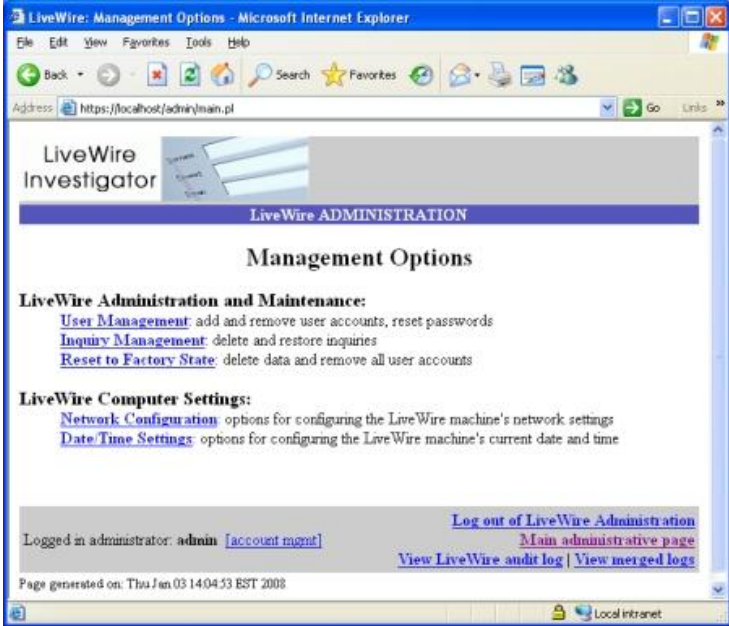
LiveWire Initial Setup, continued

Procedure: LiveWire Setup, continued

Step	Action
7	<p>Once a successful password has been set, a success verification page will be displayed. Click Continue to LiveWire Administration.</p> 

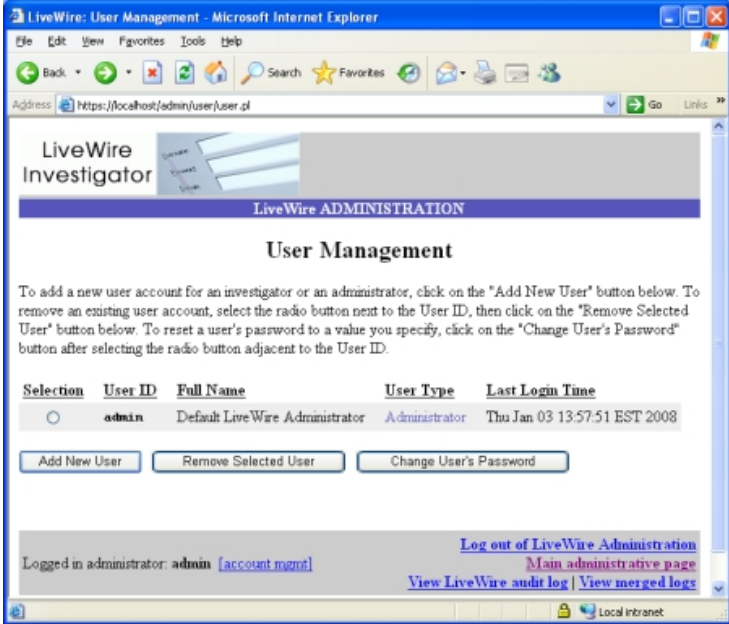
LiveWire Initial Setup, continued

Procedure: LiveWire Setup, continued

Step	Action
8	<p>The Management Options page will be displayed. Click “User Management” to continue.</p> 

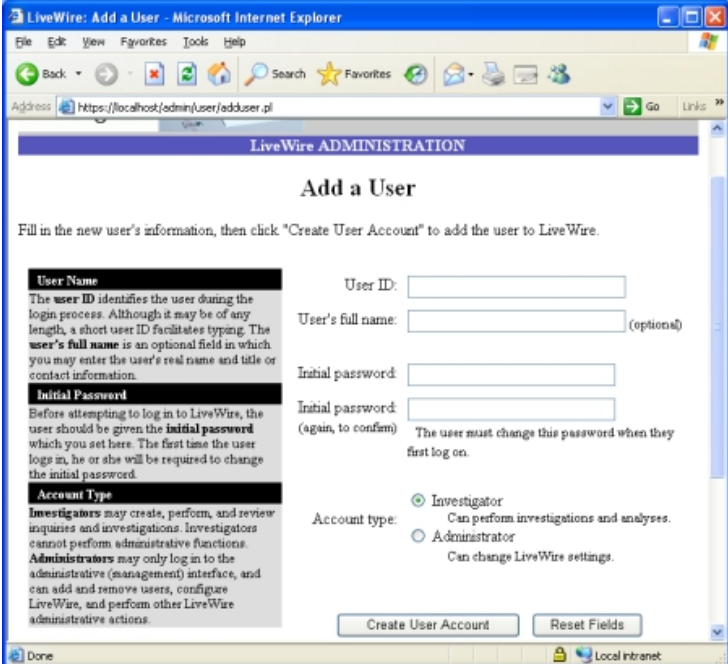
LiveWire Initial Setup, continued

Procedure: LiveWire Setup, continued

Step	Action
9	<p>The User Management page can be used to Add/Remove users or reset passwords. Click “Add New User”.</p> 

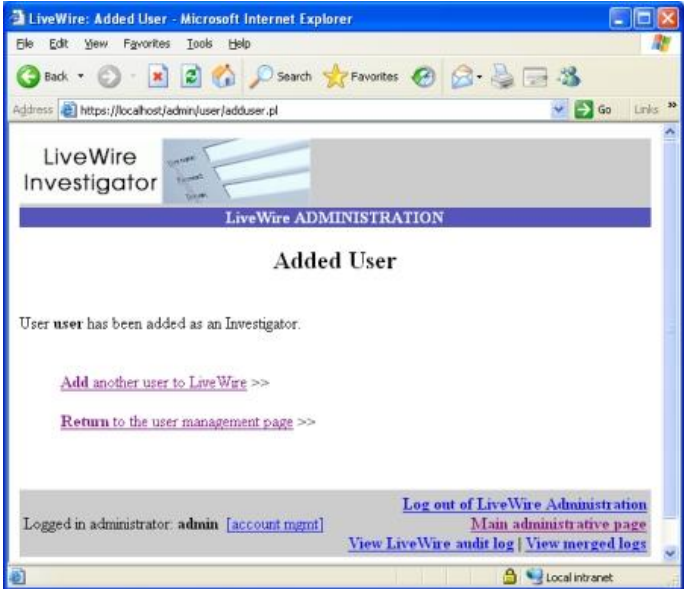
LiveWire Initial Setup, continued

Procedure: LiveWire Setup, continued

Step	Action
10	<p>To be able to use LiveWire for an investigation, there must be at least one investigator account. Fill in the appropriate user information, verify the account type as “Investigator” and click “Create User Account.”</p> 

LiveWire Initial Setup, continued

Procedure: LiveWire Setup, continued

Step	Action
11	<p>Now the user that was created can log into LiveWire and begin performing investigations.</p> 

This page intentionally left blank.

Lesson 3 – LiveDiscover

Introduction

In live network investigations, it is important to be able to effectively scan and identify the network for devices. LiveDiscover is a tool that you can use to do this.

Purpose of this Lesson

The purpose of this lesson is to introduce you to the LiveDiscover tools.

Objectives

After successfully completing this lesson, you will be able to:

- Discuss important functions of LiveDiscover
- Scan a network
- Identify devices found on the network

In this Lesson

The following table shows the contents of this lesson.

Topic	See Page
LiveDiscover Network Scanning	12-32

LiveDiscover Network Scanning

Introduction


LiveDiscover is a tool used to rapidly identify and assess resources on the network. With the information you gather from LiveDiscover, you can use LiveWire tools to perform a live analysis of a machine across the network.

LiveDiscover is able to quickly scan a range of IP addresses. You can perform an in-depth scan that reports the vulnerabilities systems. Each individual scan is stored and saved in its own database. These scans can be single targets or a whole range of IP addresses.


All information contained in LiveDiscover is stored within a database. When LiveDiscover is first started, the user can open an existing database or create a new one. To create a new database, the user simply has to provide a name that does not already exist.

LiveDiscover Interface

LiveDiscover provides a tabbed interface for navigation. The primary tabs are displayed horizontally across the top of the page. Each page displays information or options pertaining to the specific details or configurations.



The Discover tab is where the investigator enters the addresses he plans to scan. Up to four different network ranges can be scanned at the same time.



In the Network tab, a tree structure will be created showing the different devices that were discovered for that subnet as well as detailed information gathered about each of those devices.



The Responses tab displays the discovered data grouped together. Selecting any of the options will display all the items found matching that criteria.


LiveDiscover Network Scanning, continued

LiveDiscover Interface, continued




The Reports tab allows for the generation of many different report display formats. Individual types of information can be viewed in several ways. Reports can contain text as well as colored graphs.

Windows		
IP Address	NAME	OS
10.15.4.121	ODP-SALES1	Windows
10.15.4.156	ODP-TECH1	Windows
10.15.4.200	LIVEWIRE	Windows XP
10.15.4.210	SPIDER	Windows XP



The Script tab is where all the different pre-built discovery scripts are stored in the database. Customized scripts can be added to the database at any time.



The Settings tab contains the configurations to use during the live discovery process. The username and password used in this tab should be an account that has administrative rights to that device. You can also configure the SMTP setting and e-mail address to automatically have an e-mail sent once the scan is completed.



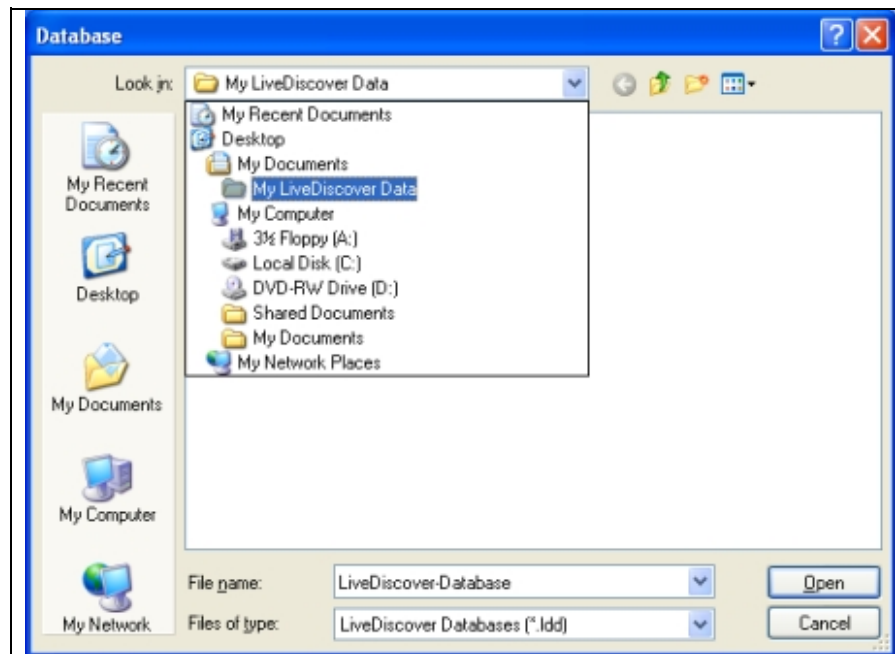
The Utilities tab offers options for scripts and results from other scans to be imported into the current database. Scripts can also be exported to be used in other scans.

LiveDiscover Network Scanning, continued

Procedure: Performing a LiveDiscover Network Scan

In this section, you will scan a subnet to determine information needed to assist with the LiveWire investigation.

Step	Action
1	Open the LiveDiscover application by navigating to: Start > All Programs > LiveDiscover > LiveDiscover.
2	You will be prompted to open a current database, or create a new database by typing in a file name. Type a name for the new database, such as "LiveDiscoverDatabase." Then click Open .



LiveDiscover Network Scanning, continued

Procedure: Performing a LiveDiscover Network Scan, continued

Step	Action
3	Because the new database name does not already exist, the user will be prompted to verify the creation of a new database. Select Yes to create the new database.



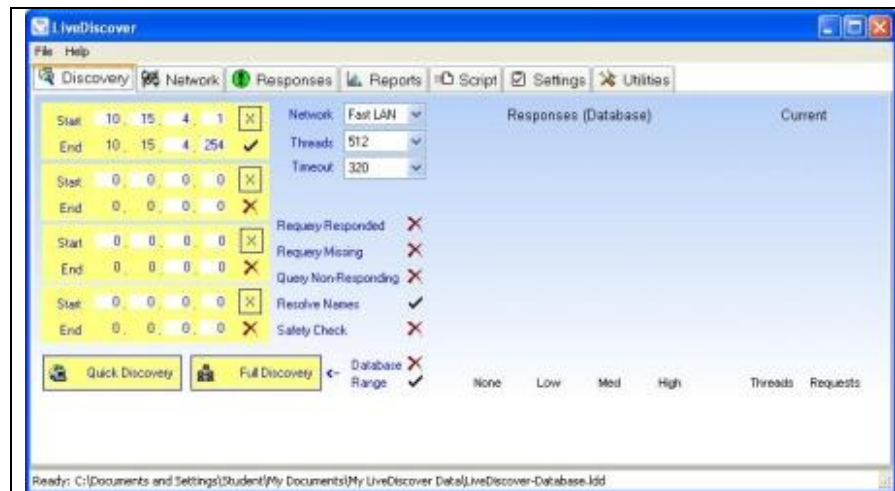
4	LiveDiscover comes with many scripts to identify devices and resources on the network. Highlight “01 – Full Discovery.scip” and click OK .
---	---



LiveDiscover Network Scanning, continued

Procedure: Performing a LiveDiscover Network Scan, continued

Step	Action
5	<p>The standard LiveDiscover page is displayed. Change the network setting to Fast LAN. Then enter in the upper box the IP address range to scan. This information will be provided by your instructor, but it will normally be an entire network range, such as the values below:</p> <p>Start IP: 10.15.4.1 End IP: 10.15.4.254</p> <p>Beside Full Discovery, click on Range to activate scanning the desired network range.</p>

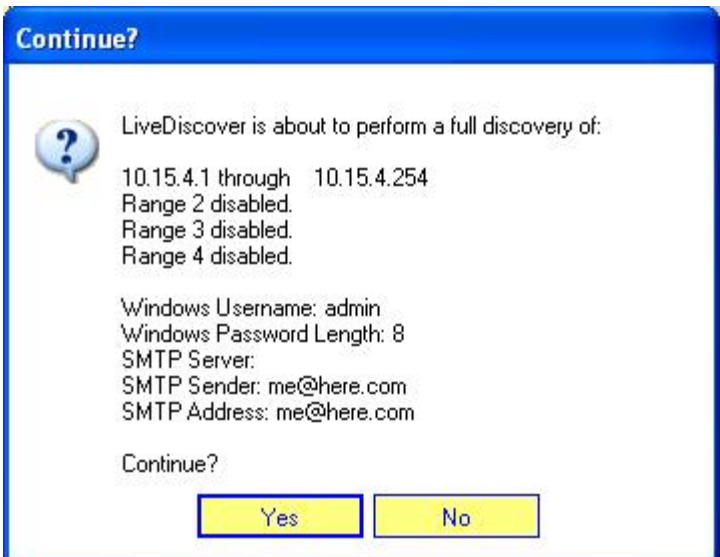


LiveDiscover Network Scanning, continued

Procedure: Performing a LiveDiscover Network Scan, continued

Step	Action
6	Click on the Settings tab and enter the following information: <User>: admin <PWD>: password

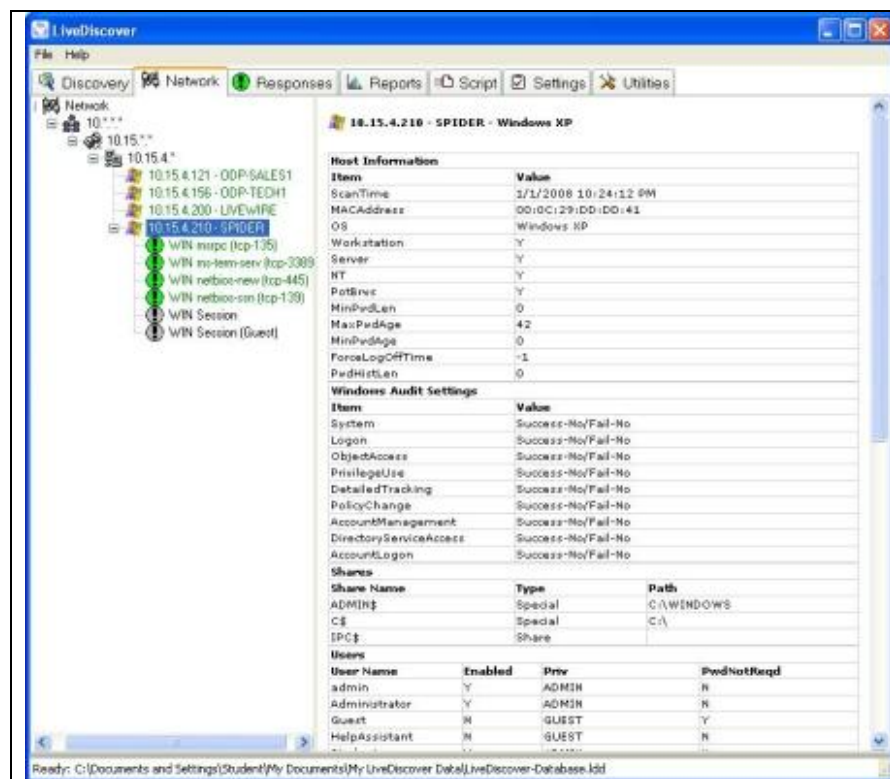


7	Go back to the Discovery tab and click Full Discovery . Click Yes to verify the information entered is correct and to begin the discovery scan. 
---	---

LiveDiscover Network Scanning, continued

Procedure: Performing a LiveDiscover Network Scan, continued

Step	Action
8	Once the scan has successfully completed, click on the Network tab. Expand the network tree until the details of the computer Spider is displayed. The information on the right verifies the operating system is a Windows XP machine.



Lesson 4 – Volatile Data Analysis

Introduction Volatile data on a system can be very valuable to an investigation. This information typically holds information regarding activity that is occurring during the investigation, but can be completely lost when the system is powered off.

Purpose of this Lesson The purpose of this lesson is to introduce you to the functionality of LiveWire Investigator.

Objectives After successfully completing this lesson, you will be able to:

- Conduct an initial inquiry of a system
- View the current open files on a system
- View the current network connections and configurations
- Image RAM over the network

In this Lesson The following table shows the contents of this lesson.

Topic	See Page
LiveWire Initial Inquiry	12-40
System State	12-49
Current User Activity	12-55
Active Network State	12-68

LiveWire Initial Inquiry

Initial Inquiry

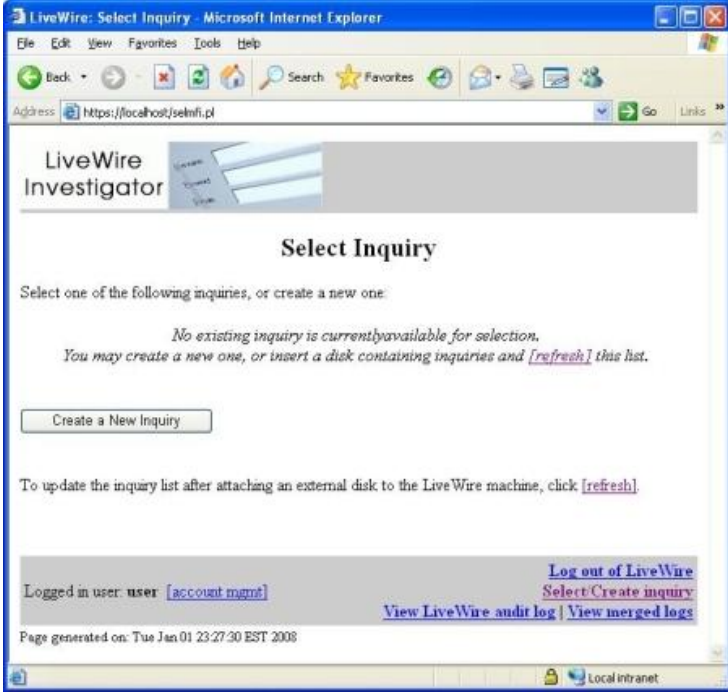
When performing an investigation or analysis of a system using LiveWire the first part of the process is the initial inquiry. This will retrieve information from the remote computer that is available at the current time. It must be noted that all live systems are dynamic; therefore, if information is gathered at a later time it may be very different. If an individual workstation is being monitored for illicit activities, the initial inquiry must be performed as those activities are occurring.

The initial inquiry and other intensive analyses do have the potential to degrade the performance of the suspect machine. Therefore, it may be possible for the user to become alerted of suspicious activity that affects his or her machine. An advanced user with authenticated credentials to the local machine may be able to determine that the analysis is taking place. These concerns must be addressed depending on the circumstances of each individual investigation.

LiveWire Initial Inquiry, continued

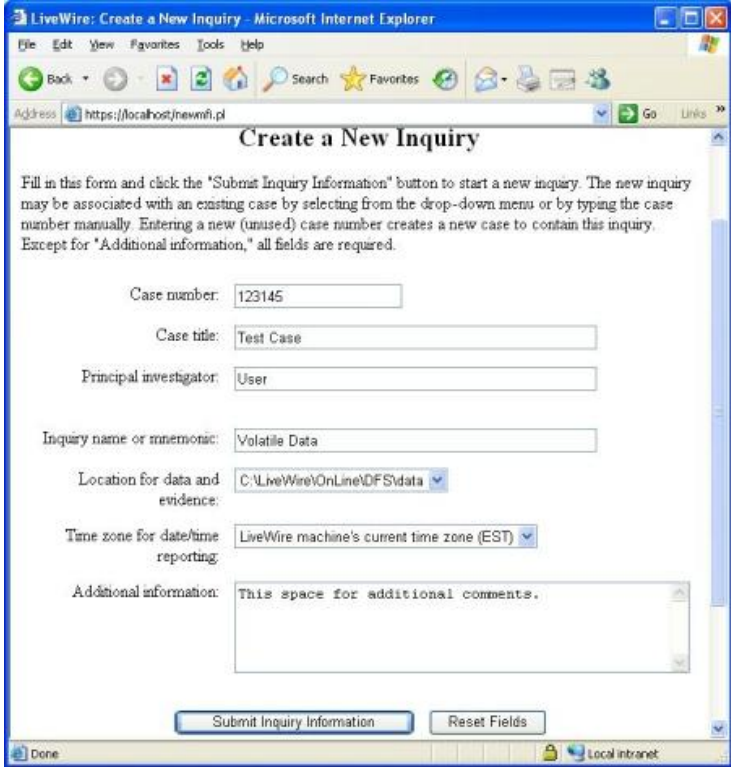
Procedure: Follow these steps to begin an initial inquiry of a Windows XP machine using LiveWire.

LiveWire – Initial Inquiry

Step	Action
1	Open LiveWire from the Windows Start Menu by selecting: “Start > All Programs > LiveWire > Launch LiveWire.” Login with the Investigator ID and password created in a previous lesson.
2	<p>To begin an investigation an Inquiry must be created. On the Select Inquiry page click Create a New Inquiry. If other inquiries already exist on the system, you would have the option to go to an existing inquiry.</p> 

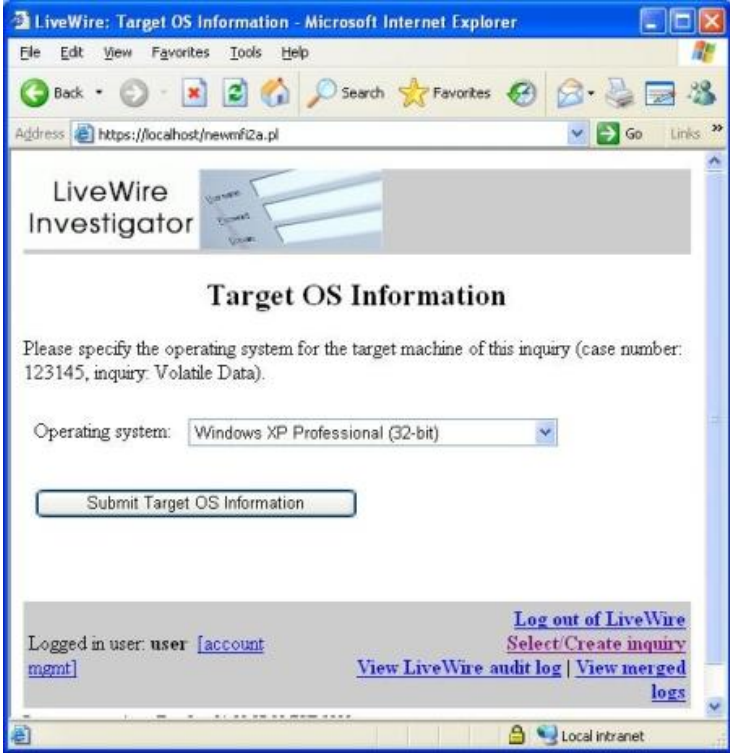
LiveWire Initial Inquiry, continued

Procedure: LiveWire – Initial Inquiry, continued

Step	Action
3	<p>Complete the inquiry information as needed, and then click Submit Inquiry Information.</p> <p><i>Notice that for these lessons we will be saving the data to the default data location. If data needs to be preserved as evidence, it's always best practice to save it to a forensically clean location, such as an external wiped drive.</i></p> 

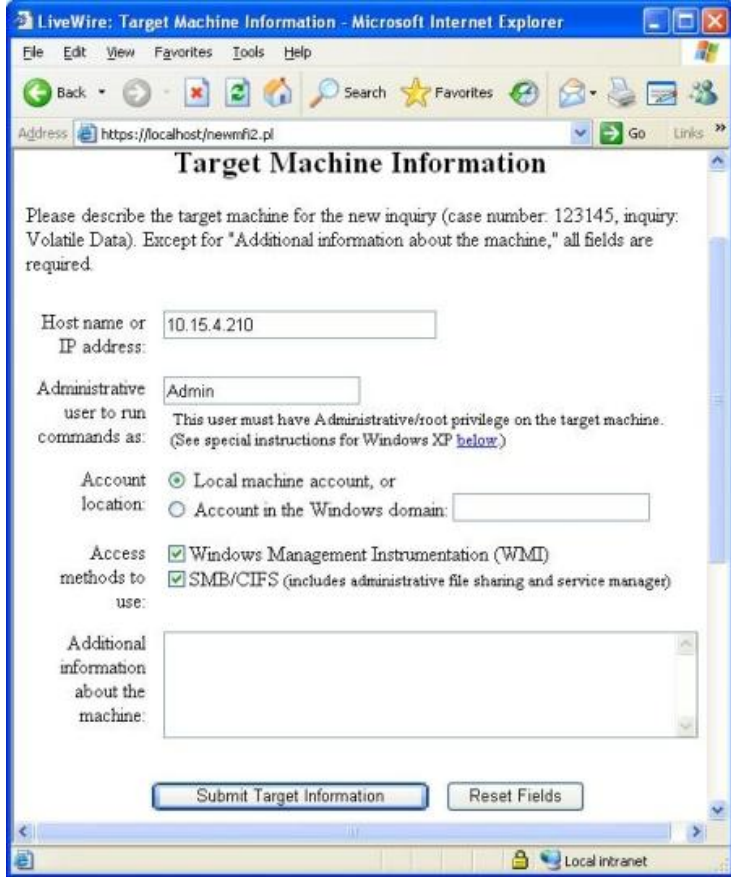
LiveWire Initial Inquiry, continued

Procedure: LiveWire – Initial Inquiry, continued

Step	Action
4	<p>Next, the target OS must be specified. Select Windows XP Professional from the drop-down list. Then click Submit Target OS Information. This information was gathered from the earlier LiveDiscover lesson.</p> 

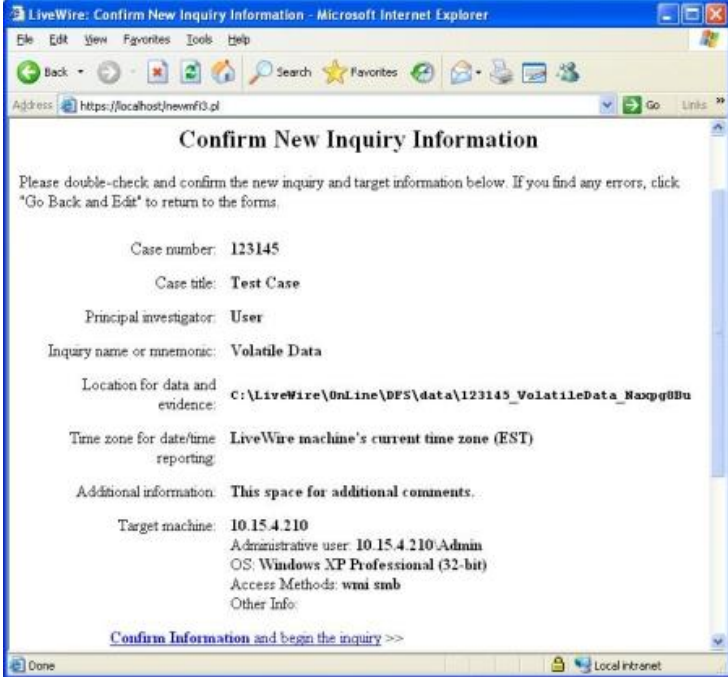
LiveWire Initial Inquiry, continued

Procedure: LiveWire – Initial Inquiry, continued

Step	Action
5	<p>The Target machine Information page is displayed. Using data gathered in the Discovery lesson, enter the following information into the boxes and click Submit Target Information.</p> <p>For classroom purposes, the IP address will be provided by your instructor.</p> <p>For the Administrative user to run commands as use: Admin</p> 

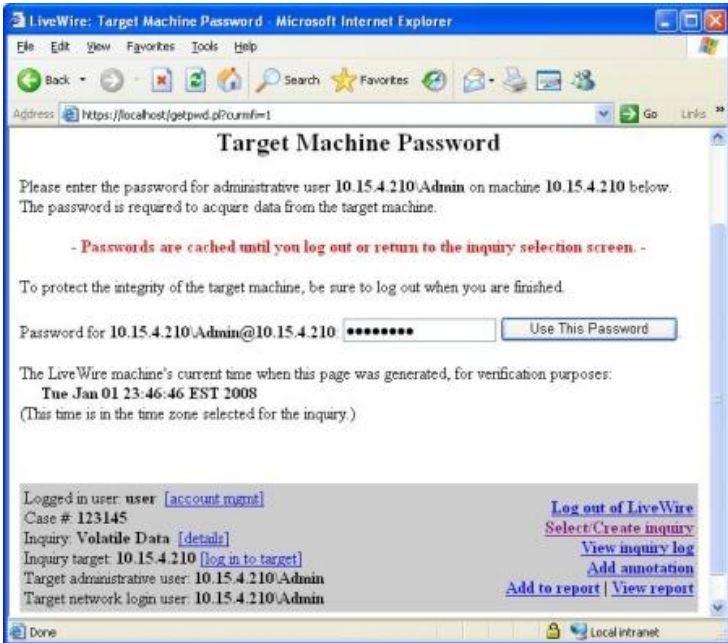
LiveWire Initial Inquiry, continued

Procedure: LiveWire – Initial Inquiry, continued

Step	Action
6	<p>The next page you to review the information entered before creating the new inquiry. Check the data and click Confirm Information and begin the inquiry to continue. Once the information has been confirmed, it can not be edited. If the information is incorrect, a new inquiry must be created with the correct information.</p> 

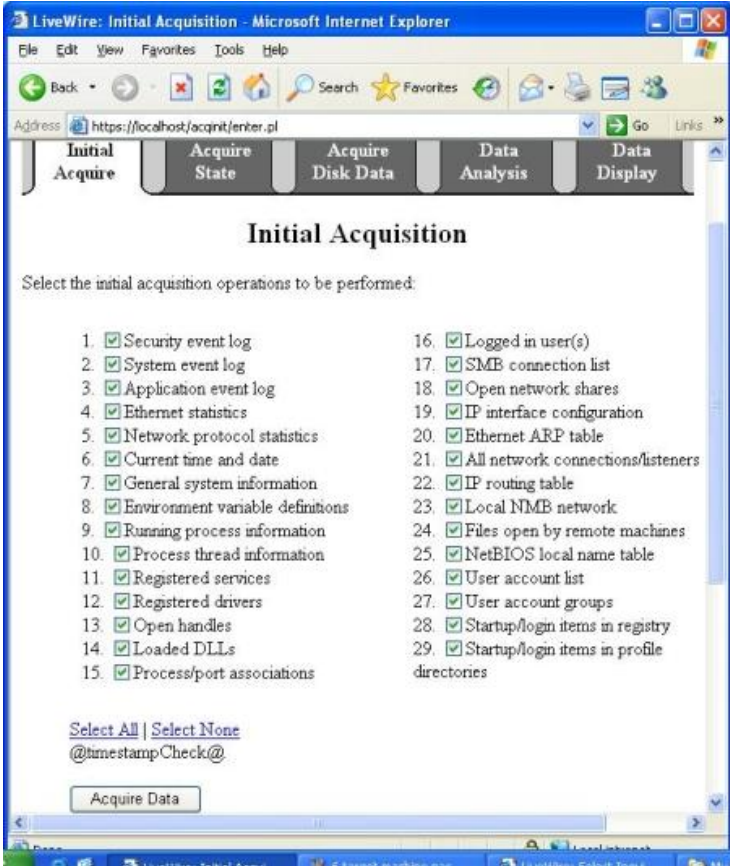
LiveWire Initial Inquiry, continued

Procedure: LiveWire – Initial Inquiry, continued

Step	Action
7	<p>Enter the password for the user on the machine that has administrative rights to the victim/suspect machine and click Use This Password.</p> <p>For classroom purposes, use the password: password</p> 

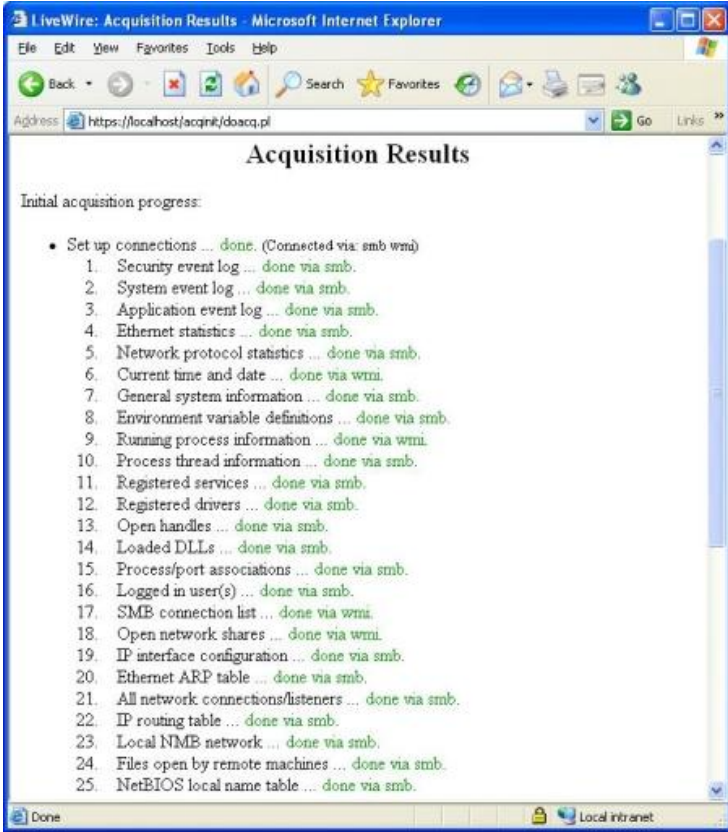
LiveWire Initial Inquiry, continued

Procedure: LiveWire – Initial Inquiry, continued

Step	Action
8	<p>The next phase is to begin the initial acquisition. This phase allows you to select all the data you wish to retrieve for analysis. Verify that all boxes are checked and click Acquire Data.</p> 

LiveWire Initial Inquiry, continued

Procedure: LiveWire – Initial Inquiry, continued

Step	Action
9	<p>Once the acquisition has completed, a page similar to the following will be displayed. This will allow you to verify that the information was retrieved from the suspect machine successfully.</p>  <p>The screenshot shows a web browser window titled "LiveWire: Acquisition Results - Microsoft Internet Explorer". The address bar displays "https://localhost/acqnit/door.pl". The page content is titled "Acquisition Results" and shows "Initial acquisition progress:". Below this, there is a bulleted list of 25 items, each followed by a status and the method used for acquisition:</p> <ul style="list-style-type: none"> • Set up connections ... done. (Connected via: smb wmi) 1. Security event log ... done via smb. 2. System event log ... done via smb. 3. Application event log ... done via smb. 4. Ethernet statistics ... done via smb. 5. Network protocol statistics ... done via smb. 6. Current time and date ... done via wmi. 7. General system information ... done via smb. 8. Environment variable definitions ... done via smb. 9. Running process information ... done via wmi. 10. Process thread information ... done via smb. 11. Registered services ... done via smb. 12. Registered drivers ... done via smb. 13. Open handles ... done via smb. 14. Loaded DLLs ... done via smb. 15. Process/port associations ... done via smb. 16. Logged in user(s) ... done via smb. 17. SMB connection list ... done via wmi. 18. Open network shares ... done via wmi. 19. IP interface configuration ... done via smb. 20. Ethernet ARP table ... done via smb. 21. All network connections/listeners ... done via smb. 22. IP routing table ... done via smb. 23. Local NMB network ... done via smb. 24. Files open by remote machines ... done via smb. 25. NetBIOS local name table ... done via smb. <p>The status bar at the bottom of the browser window shows "Done" and "Local intranet".</p>

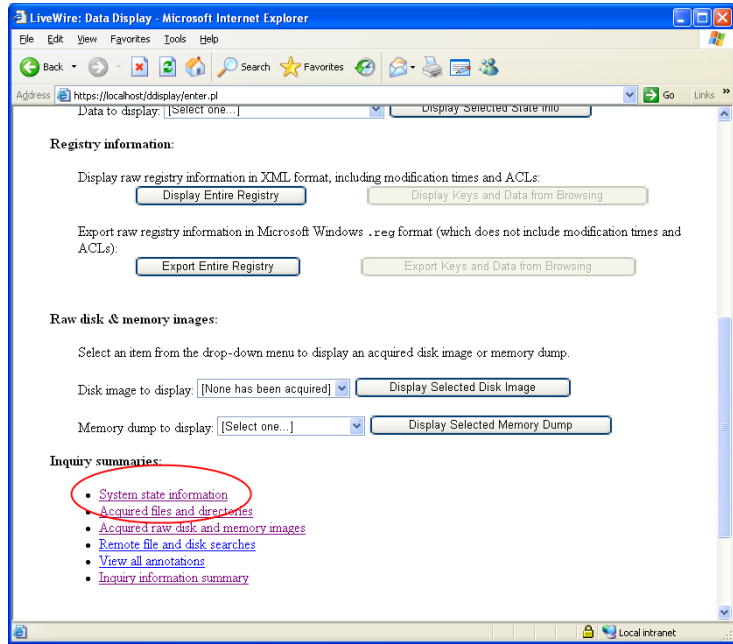
System State

Overview

The previous section explained how to conduct the LiveWire initial inquiry on the system for analysis. This section will continue to explain how to capture volatile data on a system in order for the data to be analyzed.

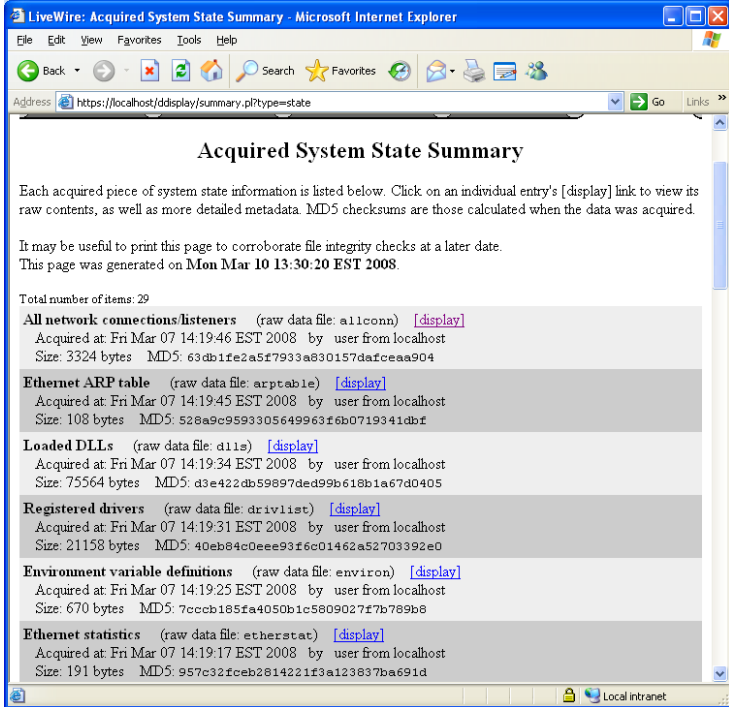
Procedure: LiveWire – Display Acquired System State Summary

Use these steps to view the acquired system state summary. This section continues from the previous section.

Step	Action
1	Click the Data Display tab.
2	<p>Scroll down the Inquiry Summaries section. Click the System state information link.</p> 

System State, continued

Procedure: LiveWire – Display Acquired System State Summary, continued

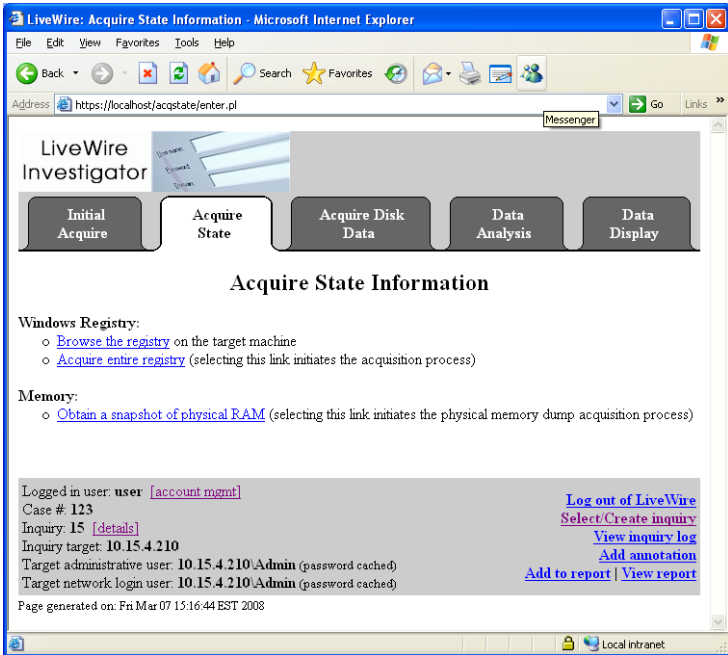
Step	Action
3	<p>The acquired system state summary page will be displayed. Links for details about each item is available for all items. You may analyze information gathered by selecting the [display] link next to each item.</p>  <p>LiveWire: Acquired System State Summary - Microsoft Internet Explorer</p> <p>File Edit View Favorites Tools Help</p> <p>Back Forward Stop Reload Home Search Favorites Print Mail News RSS Feeds</p> <p>Address https://localhost/ddisplay/summary.pl?type=state Go Links</p> <h3>Acquired System State Summary</h3> <p>Each acquired piece of system state information is listed below. Click on an individual entry's [display] link to view its raw contents, as well as more detailed metadata. MD5 checksums are those calculated when the data was acquired.</p> <p>It may be useful to print this page to corroborate file integrity checks at a later date. This page was generated on Mon Mar 10 13:30:20 EST 2008.</p> <p>Total number of items: 29</p> <p>All network connections/listeners (raw data file: allconn) [display] Acquired at: Fri Mar 07 14:19:46 EST 2008 by user from localhost Size: 3324 bytes MD5: 63db1fe2a5f7933a830157dafceaa904</p> <p>Ethernet ARP table (raw data file: arptable) [display] Acquired at: Fri Mar 07 14:19:45 EST 2008 by user from localhost Size: 108 bytes MD5: 528a9c9593305649963f6b0719341dbf</p> <p>Loaded DLLs (raw data file: dlls) [display] Acquired at: Fri Mar 07 14:19:34 EST 2008 by user from localhost Size: 75564 bytes MD5: d3e422db59897ded99b618b1a67d0405</p> <p>Registered drivers (raw data file: driv11st) [display] Acquired at: Fri Mar 07 14:19:31 EST 2008 by user from localhost Size: 21158 bytes MD5: 40eb84c0eee93f6c01462a52703392e0</p> <p>Environment variable definitions (raw data file: environ) [display] Acquired at: Fri Mar 07 14:19:25 EST 2008 by user from localhost Size: 670 bytes MD5: 7cccb185fa4050b1c5809027f7b789b8</p> <p>Ethernet statistics (raw data file: etherstat) [display] Acquired at: Fri Mar 07 14:19:17 EST 2008 by user from localhost Size: 191 bytes MD5: 957c32fceb2814221f3a123837ba691d</p> <p>Local intranet</p>

System State, continued

Procedure: LiveWire – Acquire Physical RAM and the Registry

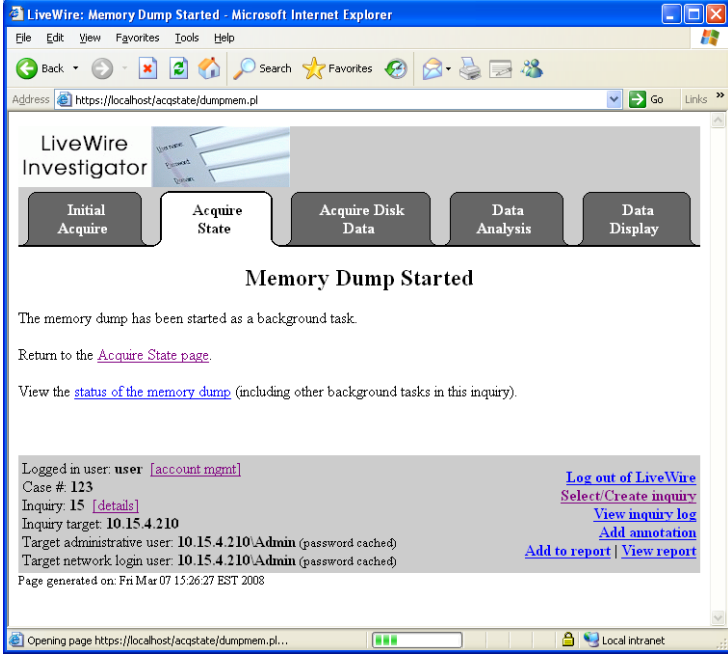
Follow these steps to use LiveWire Investigator to acquire a snapshot of the RAM and acquire entire system registry. It is always best practice to acquire the most volatile data prior to collecting other less volatile types of data.

This section continues from the previous section.

Step	Action
1	<p>In the tab area at the top, click the Acquire State tab.</p> 

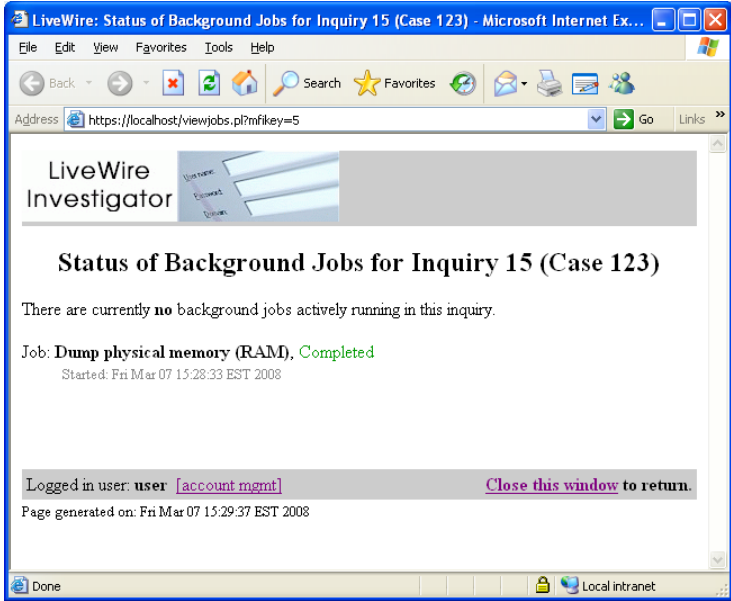
System State, continued

Procedure: LiveWire – Acquire Registry and Physical RAM, continued

Step	Action
2	<p>Click Obtain a snapshot of physical RAM. This may take a few minutes depending on the size of RAM and current load on the system.</p> 

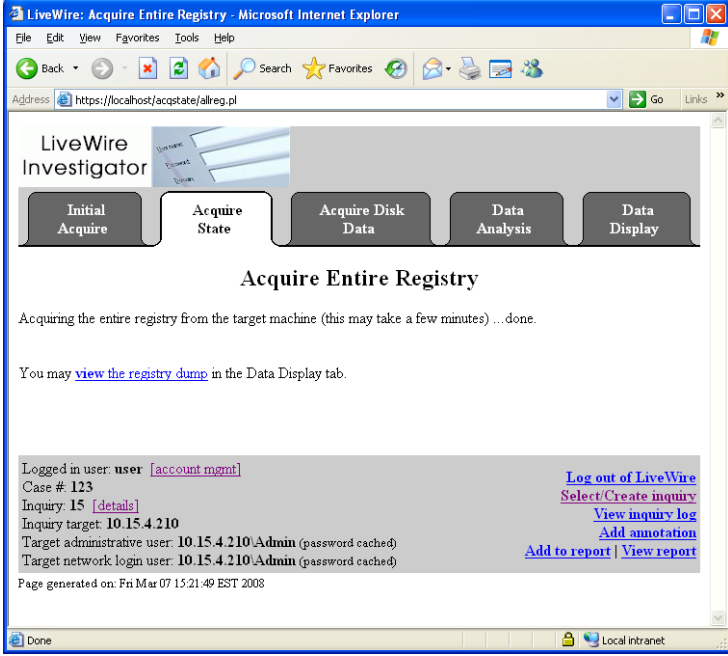
System State, continued

Procedure: LiveWire – Acquire Registry and Physical RAM, continued

Step	Action
4	<p>Because it may take a few minutes to complete, the Investigator can click Status of the memory dump to view the physical memory dump status as it is being captured.</p> 

System State, continued

Procedure: LiveWire – Acquire Registry and Physical RAM, continued

Step	Action
4	<p>Click the Acquire Tab, and then click Acquire Entire Registry. This process may take a few minutes to complete.</p> 

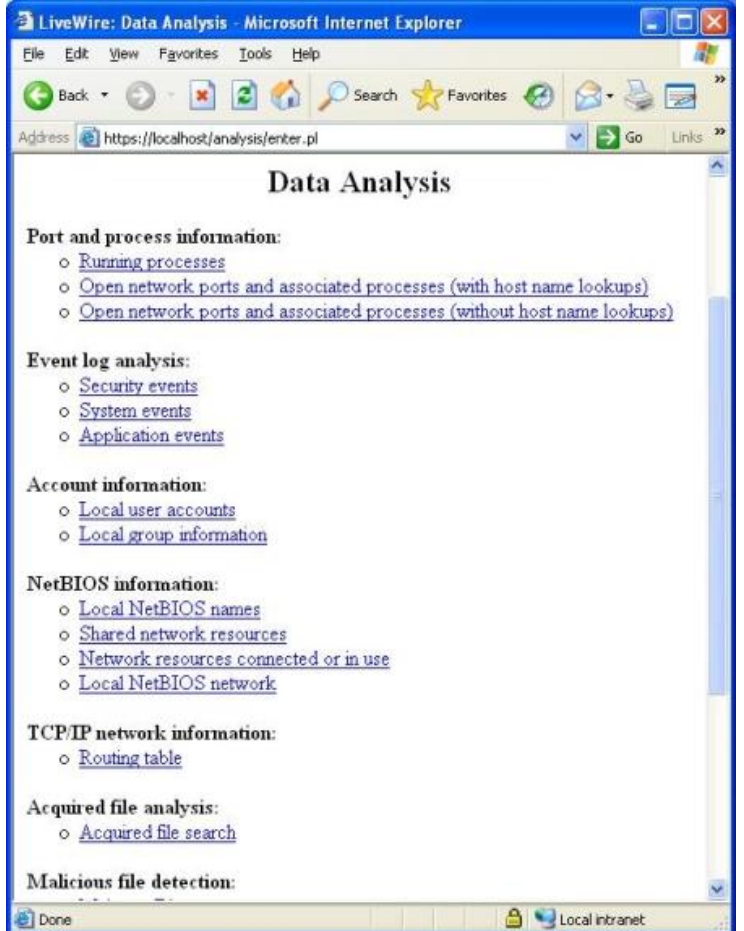
Current User Activity

Overview

The previous section explained how to gather the current information about the system. This section will continue to explain how to examine the gathered data to see what the current user's activities are by looking at what is currently open and running on the system.

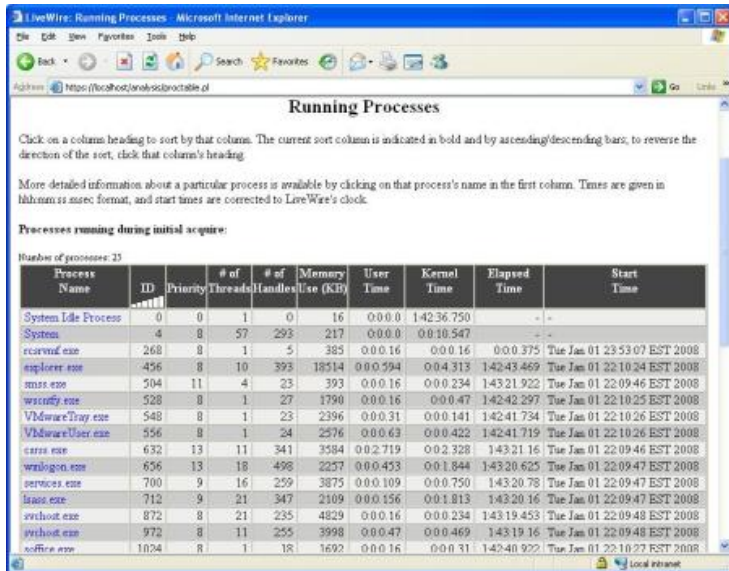
Procedure: LiveWire – Current User Activity

Follow these steps to conduct an analysis of the current user activities on the system using LiveWire. This section continues from the previous section.

Step	Action
1	<p>Click on the Data Analysis tab to display the following page. This page contains links to view the various types of information previously gathered.</p> 

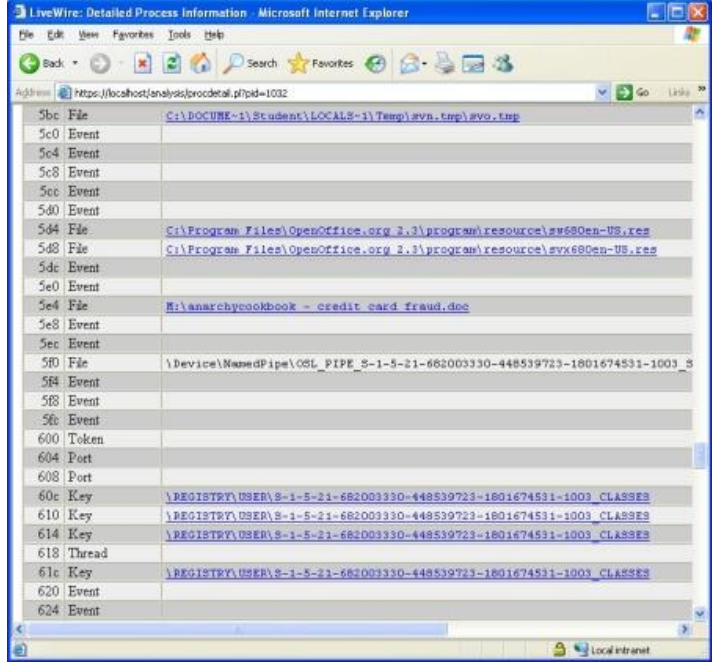
Current User Activity, continued

Procedure: LiveWire – Current User Activity, continued

Step	Action
2	<p>Click the Running Processes link to view the active process currently running on the machine. Notice that some of the process running is TrueCrypt.exe, soffice.exe and soffice.bin. This tells us a little bit about what the current user is doing on the system at this time.</p> 
3	<p>All the items under the Process Name are currently running on the system. Click on soffice.bin to see what is associated with that active process.</p>

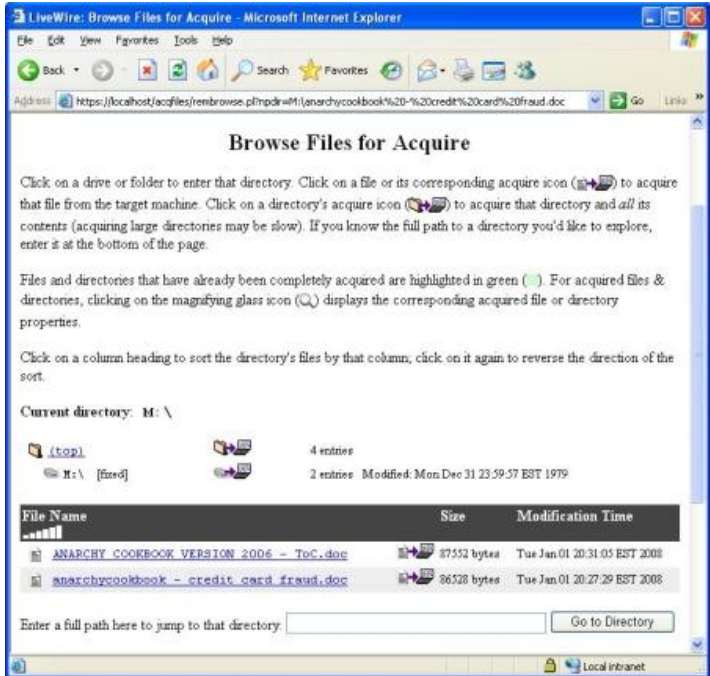
Current User Activity, continued

Procedure: LiveWire – Current User Activity, continued

Step	Action
4	<p>Scroll down to see what .doc file is currently opened by the soffice.bin process. Click on the link M:\anarchycookbook – credit card fraud.doc.</p> 


Current User Activity, continued

Procedure: LiveWire – Current User Activity, continued

Step	Action
5	<p>The location of that file opens and you can view details about the file. This also allows you to view what other files are stored in that directory. Click on anarchycookbook – credit card fraud.doc to acquire the file.</p> 

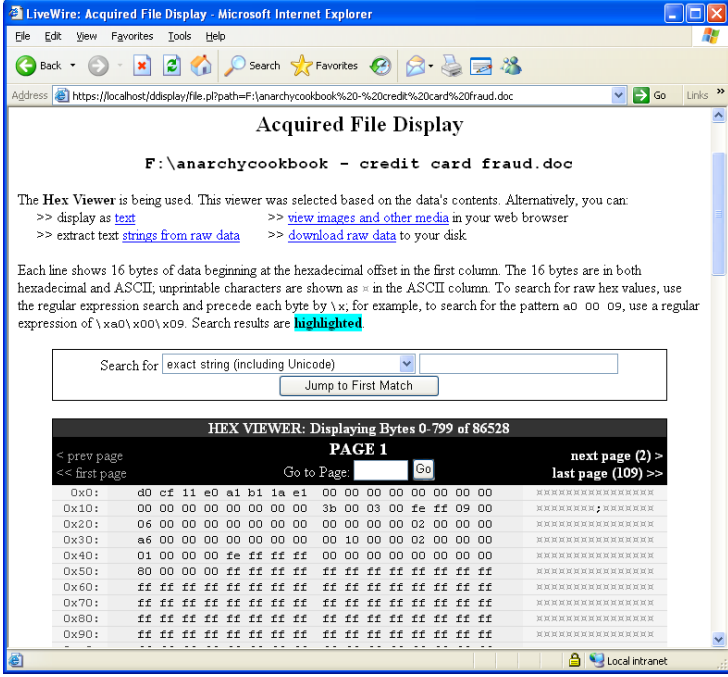
Current User Activity, continued

Procedure: LiveWire – Current User Activity, continued

Step	Action
6	<p>Now the file has been successfully acquired. To view the file in the hex viewer, click view the acquired file.</p> 

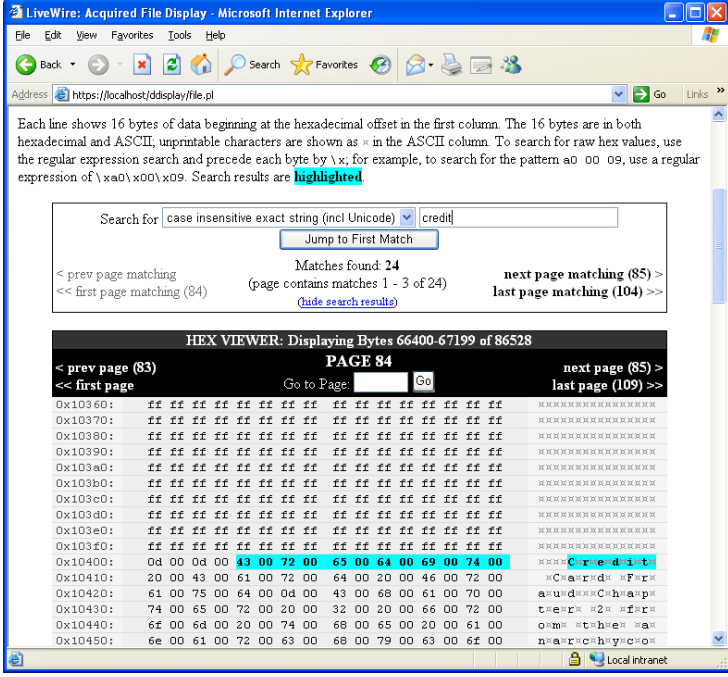
Current User Activity, continued

Procedure: LiveWire – Current User Activity, continued

Step	Action
7	<p>The Hex Viewer is the default view mode. You can search by keywords or browse around the file in this mode.</p> 

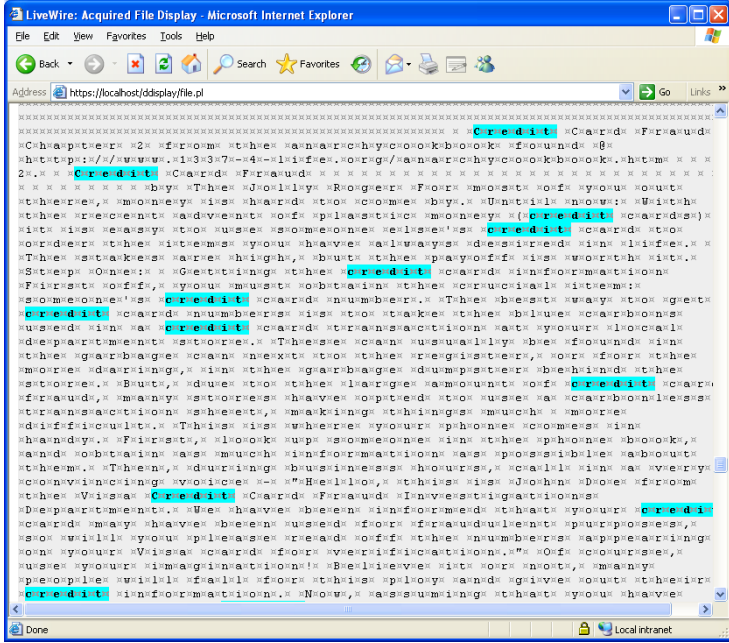
Current User Activity, continued

Procedure: LiveWire – Current User Activity, continued

Step	Action
8	<p>In the Search for pull-down menu, select case insensitive exact string (include Unicode). Then input the word credit as the string. Click Jump to First Match.</p> <p>How many matches were found? The first match was on what page?</p> 

Current User Activity, continued

Procedure: LiveWire – Current User Activity, continued

Step	Action
9	<p>Now let's change over to another useful view mode, the text view mode. At the top of the page, click the link text beside display as.</p> 

Current User Activity, continued

Procedure:**Viewing Acquired Files on the Local File System**

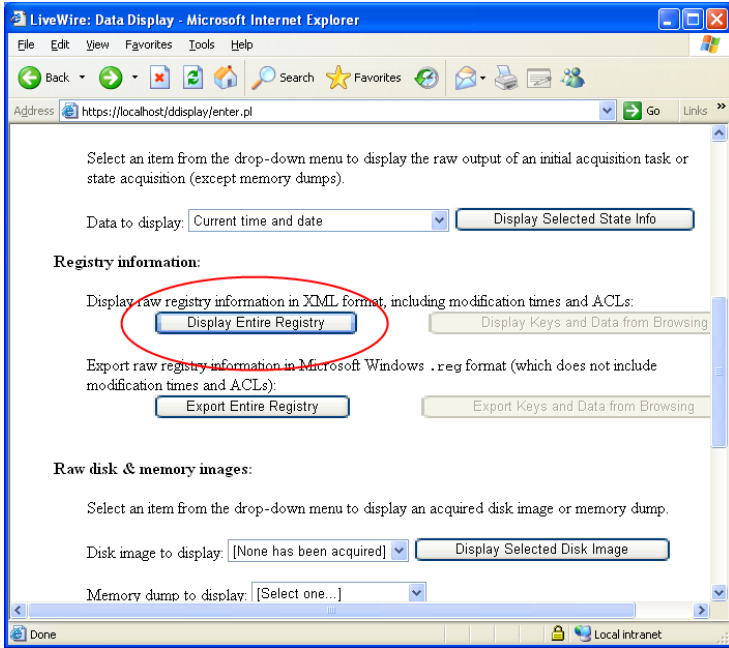
LiveWire only supports viewing image type files that are acquired. The acquired files do not retain the file extension in the local file system. The following procedure is used to rename acquired files in order to view them with their native application.

Step	Action
1	Open Windows Explorer and browse to the data directory at: C:\LiveWire\OnLine\DFS\data. This folder is where data for each case is stored.
2	Each folder under the data directory is a different case. Directory names contain the Case Number then Inquiry Name followed by random characters. Select the current case folder for the previous lessons and open it. Example: 12345_VolatileData_Naxpg8Bu
3	Now navigate into the rawdata\files folder.
4	Search through the randomly named folders for the file that was previously downloaded.
5	Rename the file with the .doc extension. Note: Renaming a file does <i>not</i> alter its hash value.
6	Double click the file to open it in MS Word or WordPad.

Current User Activity, continued

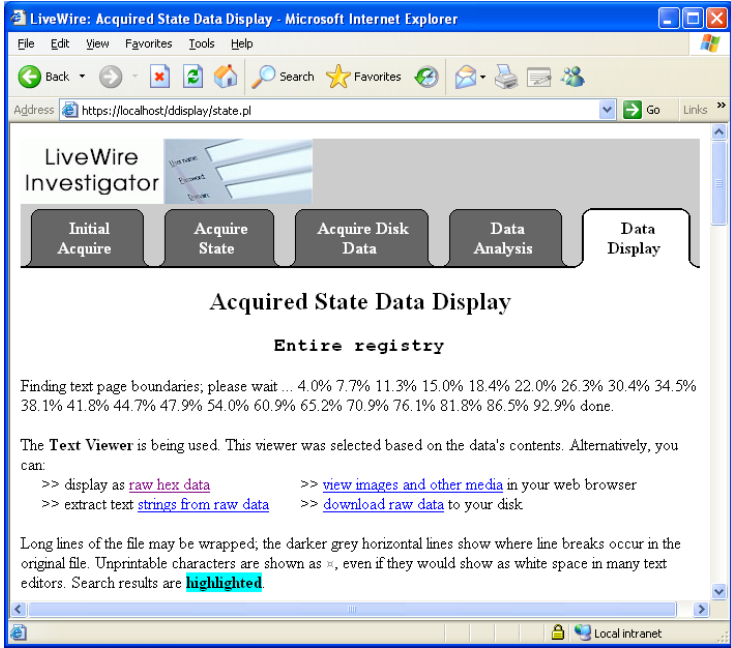
Procedure: LiveWire – Display Registry Information

Follow these steps to conduct an analysis of the current user activities on the system using LiveWire. This section continues from the previous section.

Step	Action
1	<p>The registry captured in the previous section can also be examined. Click on Data Display tab, scroll down and click on Display Entire Registry.</p> 

Current User Activity, continued

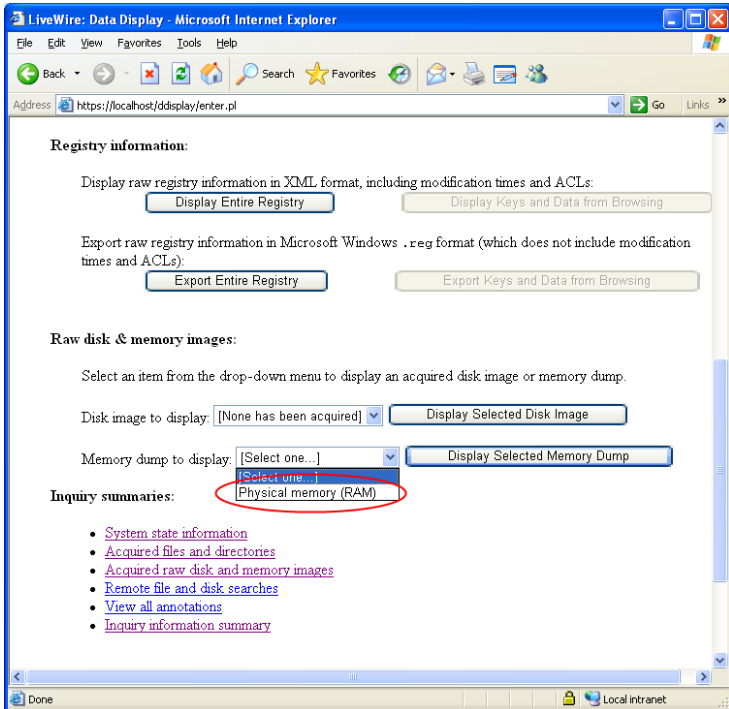
Procedure: LiveWire – Display Registry Information, continued

Step	Action
2	<p>This process may take some time. A screen similar to the one below will be displayed once completed. The registry can be searched the same way as previous exercises.</p> 

Current User Activity, continued

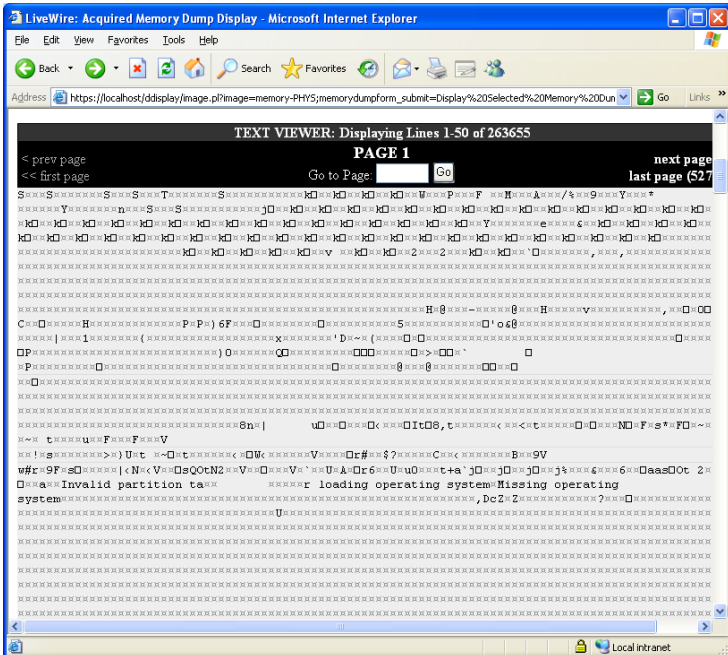
Procedure: Display Physical RAM Dump

The physical RAM dump can contain some of the most critical evidence in any case. This section will explain how to open and view the physical RAM dump in LiveWire Investigator.

Step	Action
1	<p>Click on Data Display tab, scroll down to the Memory dump to display pull-down menu and select Physical memory (RAM). Click on Display Selected memory Dump.</p> 

Current User Activity, continued

Procedure: Display Physical RAM Dump, continued

Step	Action
2	<p>The physical RAM dump can be viewed and searched like any other piece of evidence.</p> 

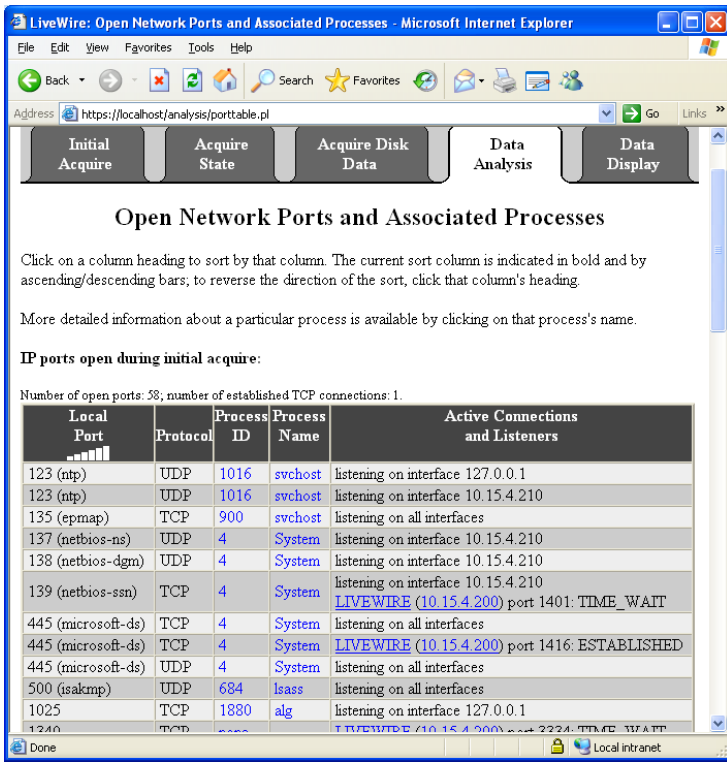
Active Network State

Overview

LiveWire Investigator has the ability to save and display the current state of the network connections and configurations.

Procedure: Display Captured Network Details

This section will explain how to view the different options available for viewing network details.

Step	Action																																																																	
1	<p>Click on the Data Analysis tab. Then click Open network ports and associated processes. It should be known that ports 3334 and 3335 will be included due to LiveWire connecting to the system to retrieve data.</p>  <p>Open Network Ports and Associated Processes</p> <p>Click on a column heading to sort by that column. The current sort column is indicated in bold and by ascending/descending bars; to reverse the direction of the sort, click that column's heading.</p> <p>More detailed information about a particular process is available by clicking on that process's name.</p> <p>IP ports open during initial acquire:</p> <p>Number of open ports: 38; number of established TCP connections: 1.</p> <table><tr><th>Local Port</th><th>Protocol</th><th>Process ID</th><th>Process Name</th><th>Active Connections and Listeners</th></tr><tr><td>123 (ntp)</td><td>UDP</td><td>1016</td><td>svchost</td><td>listening on interface 127.0.0.1</td></tr><tr><td>123 (ntp)</td><td>UDP</td><td>1016</td><td>svchost</td><td>listening on interface 10.15.4.210</td></tr><tr><td>135 (epmap)</td><td>TCP</td><td>900</td><td>svchost</td><td>listening on all interfaces</td></tr><tr><td>137 (netbios-ns)</td><td>UDP</td><td>4</td><td>System</td><td>listening on interface 10.15.4.210</td></tr><tr><td>138 (netbios-dgm)</td><td>UDP</td><td>4</td><td>System</td><td>listening on interface 10.15.4.210</td></tr><tr><td>139 (netbios-ssn)</td><td>TCP</td><td>4</td><td>System</td><td>listening on interface 10.15.4.210 LIVEWIRE (10.15.4.200) port 1401: TIME_WAIT</td></tr><tr><td>445 (microsoft-ds)</td><td>TCP</td><td>4</td><td>System</td><td>listening on all interfaces</td></tr><tr><td>445 (microsoft-ds)</td><td>TCP</td><td>4</td><td>System</td><td>LIVEWIRE (10.15.4.200) port 1416: ESTABLISHED</td></tr><tr><td>445 (microsoft-ds)</td><td>UDP</td><td>4</td><td>System</td><td>listening on all interfaces</td></tr><tr><td>500 (isakmp)</td><td>UDP</td><td>684</td><td>lsass</td><td>listening on all interfaces</td></tr><tr><td>1025</td><td>TCP</td><td>1880</td><td>alg</td><td>listening on interface 127.0.0.1</td></tr><tr><td>1240</td><td>TCP</td><td>...</td><td>...</td><td>LIVEWIRE (10.15.4.200) port 2224: TIME_WAIT</td></tr></table>	Local Port	Protocol	Process ID	Process Name	Active Connections and Listeners	123 (ntp)	UDP	1016	svchost	listening on interface 127.0.0.1	123 (ntp)	UDP	1016	svchost	listening on interface 10.15.4.210	135 (epmap)	TCP	900	svchost	listening on all interfaces	137 (netbios-ns)	UDP	4	System	listening on interface 10.15.4.210	138 (netbios-dgm)	UDP	4	System	listening on interface 10.15.4.210	139 (netbios-ssn)	TCP	4	System	listening on interface 10.15.4.210 LIVEWIRE (10.15.4.200) port 1401: TIME_WAIT	445 (microsoft-ds)	TCP	4	System	listening on all interfaces	445 (microsoft-ds)	TCP	4	System	LIVEWIRE (10.15.4.200) port 1416: ESTABLISHED	445 (microsoft-ds)	UDP	4	System	listening on all interfaces	500 (isakmp)	UDP	684	lsass	listening on all interfaces	1025	TCP	1880	alg	listening on interface 127.0.0.1	1240	TCP	LIVEWIRE (10.15.4.200) port 2224: TIME_WAIT
Local Port	Protocol	Process ID	Process Name	Active Connections and Listeners																																																														
123 (ntp)	UDP	1016	svchost	listening on interface 127.0.0.1																																																														
123 (ntp)	UDP	1016	svchost	listening on interface 10.15.4.210																																																														
135 (epmap)	TCP	900	svchost	listening on all interfaces																																																														
137 (netbios-ns)	UDP	4	System	listening on interface 10.15.4.210																																																														
138 (netbios-dgm)	UDP	4	System	listening on interface 10.15.4.210																																																														
139 (netbios-ssn)	TCP	4	System	listening on interface 10.15.4.210 LIVEWIRE (10.15.4.200) port 1401: TIME_WAIT																																																														
445 (microsoft-ds)	TCP	4	System	listening on all interfaces																																																														
445 (microsoft-ds)	TCP	4	System	LIVEWIRE (10.15.4.200) port 1416: ESTABLISHED																																																														
445 (microsoft-ds)	UDP	4	System	listening on all interfaces																																																														
500 (isakmp)	UDP	684	lsass	listening on all interfaces																																																														
1025	TCP	1880	alg	listening on interface 127.0.0.1																																																														
1240	TCP	LIVEWIRE (10.15.4.200) port 2224: TIME_WAIT																																																														

Active Network State, continued

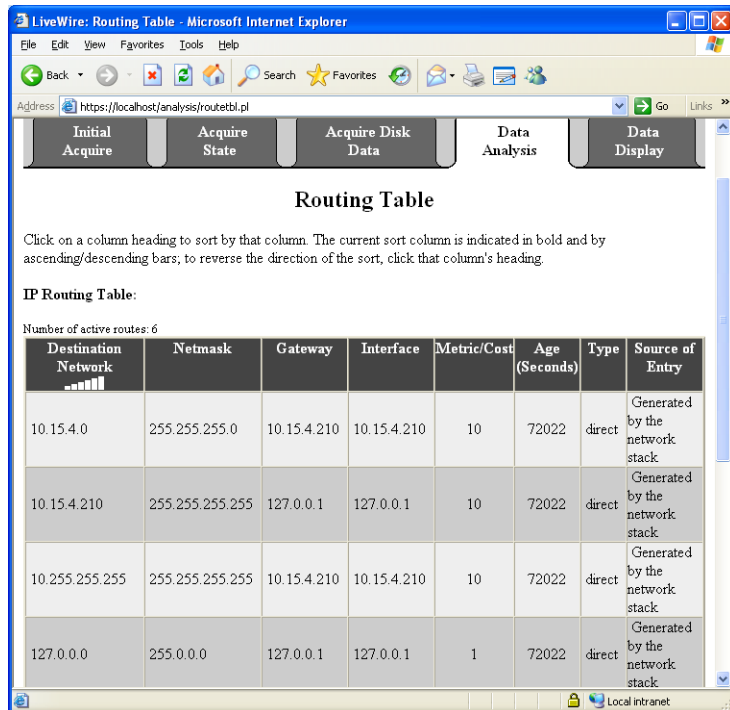
Procedure: Display Captured Network Details, continued

Step

2

Action

Click on the **Data Analysis** tab. Then click **Routing table**.



LiveWire: Routing Table - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Back Forward Stop Reload Home Search Favorites Local intranet

Address: <https://localhost/analysis/routetbl.pl> Go Links

Initial Acquire Acquire State Acquire Disk Data **Data Analysis** Data Display

Routing Table

Click on a column heading to sort by that column. The current sort column is indicated in bold and by ascending/descending bars; to reverse the direction of the sort, click that column's heading.

IP Routing Table:

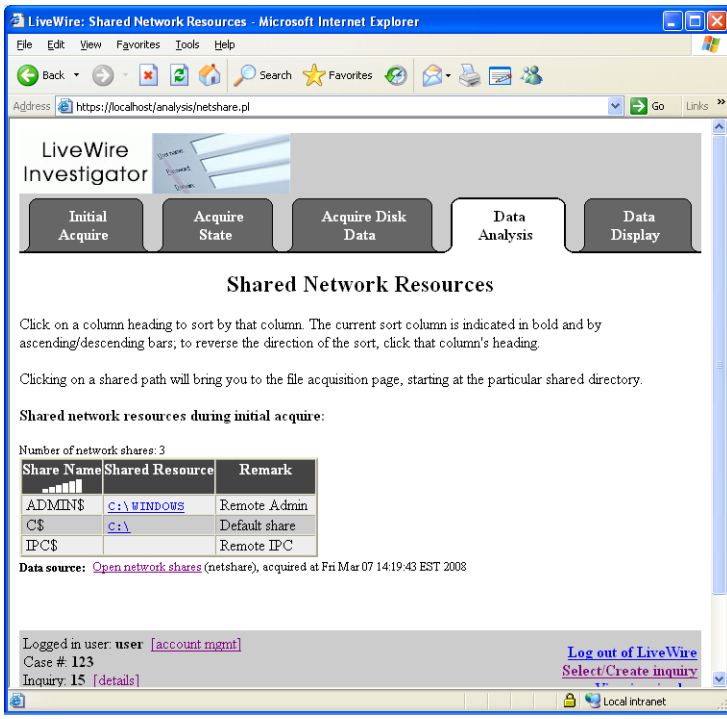
Number of active routes: 6

Destination Network	Netmask	Gateway	Interface	Metric/Cost	Age (Seconds)	Type	Source of Entry
10.15.4.0	255.255.255.0	10.15.4.210	10.15.4.210	10	72022	direct	Generated by the network stack
10.15.4.210	255.255.255.255	127.0.0.1	127.0.0.1	10	72022	direct	Generated by the network stack
10.255.255.255	255.255.255.255	10.15.4.210	10.15.4.210	10	72022	direct	Generated by the network stack
127.0.0.0	255.0.0.0	127.0.0.1	127.0.0.1	1	72022	direct	Generated by the network stack

Local intranet

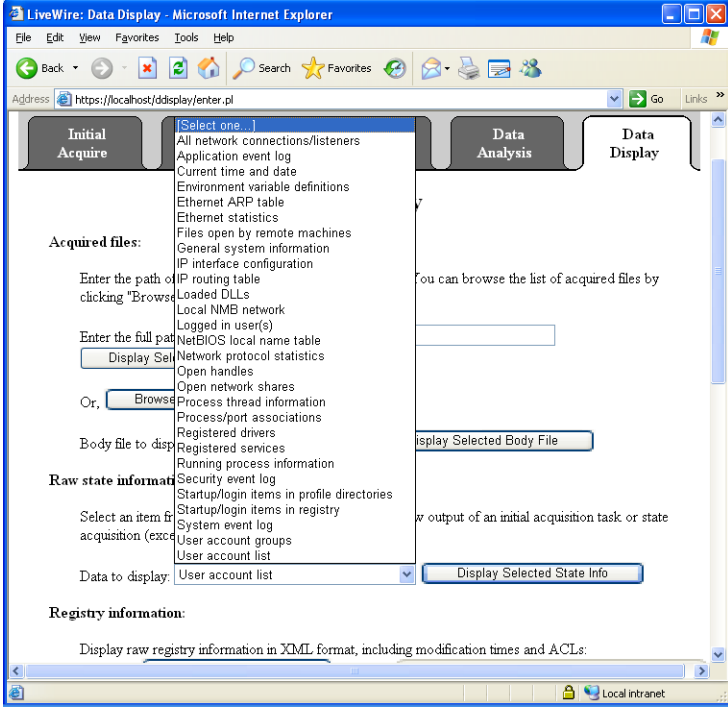
Active Network State, continued

Procedure: Display Captured Network Details, continued

Step	Action
3	<p>Click on the Data Analysis tab. Then click Shared network resources.</p> 

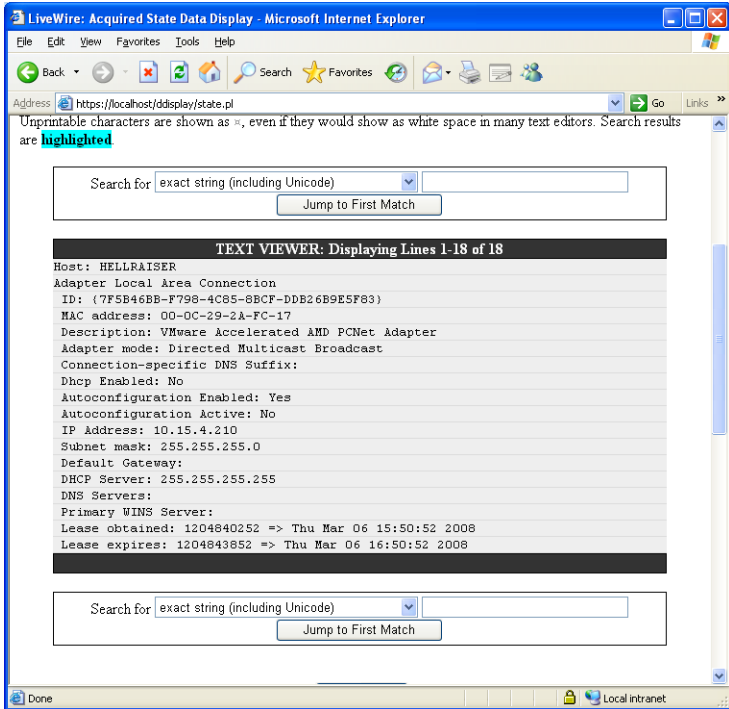
Active Network State, continued

Procedure: Display Captured Network Details, continued

Step	Action
4	<p>Another way to display network information is in the data pulled during the initial acquire. Click the Data Display tab. Under the Raw state information, in the drop-down menu, select IP interface configuration.</p> <p><i>Notice: This is another way to view the same data in other places.</i></p> <p><i>Example: Click the Data Display tab, and then click System State Information at the bottom of the page. This will display this data on a single page.</i></p> 

Active Network State, continued

Procedure: Display Captured Network Details, continued

Step	Action
5	<p>Then click Display Selected State Info to view a display similar to the one below.</p> 
6	<p>Click the Back button or Display Data tab and take a few minutes to browse around other available options, such as:</p> <ul style="list-style-type: none"> • Ethernet ARP table • Ethernet statistics • Local NMB network • NetBIOS local name table • Network protocol statistics • Open network shares

Lesson 5 – Evidence Collection

Introduction

Gathering evidence and creating disk images are critical to every investigation. LiveWire allows you to conduct an investigation of a remote system.

Purpose of this Lesson

The purpose of this lesson is to explain how to use LiveWire Investigator to collect evidence from a suspect system.

Objectives

After successfully completing this lesson, you will be able to:

- Determine the status of the file system
- Generate a disk image
- Collect file evidence from the remote system

In this Lesson

The following table shows the contents of this lesson.

Topic	See Page
File System Status	12-74
Physical vs. Logical	12-78
Collection and Preservation	12-84
Hashing	12-88

File System Status

Overview

In this lesson, you will learn how to examine physical and logical disk structure. This information can be valuable to the overall investigation and should be included in the documentation.

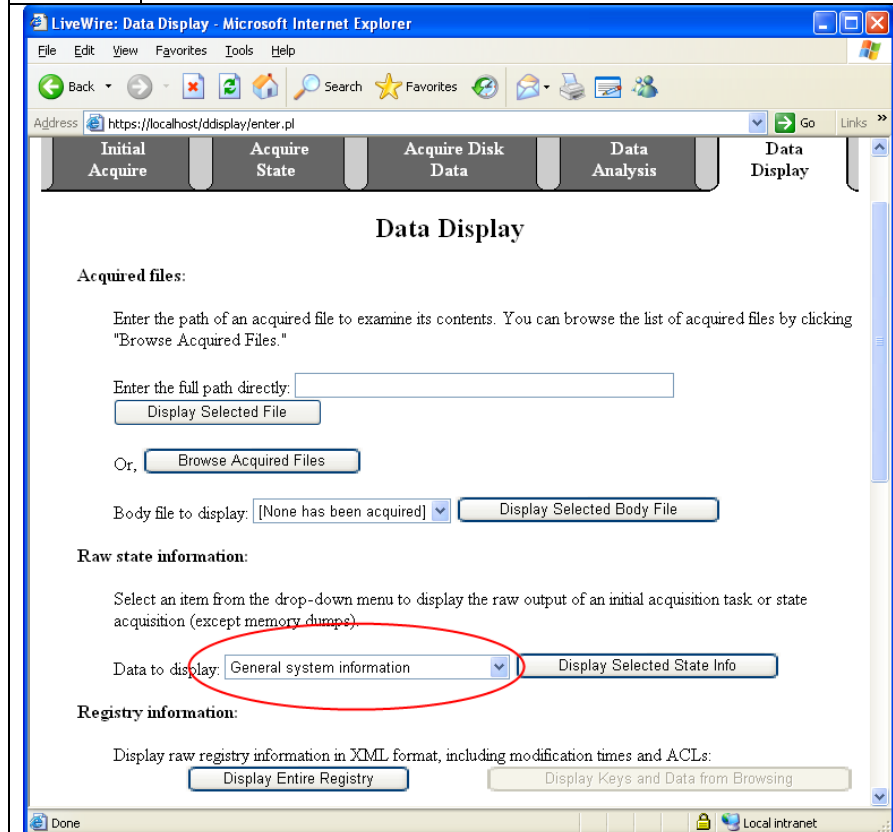
Physical Disk Size

The physical size of the suspect system is a critical part of any investigation. The investigator must have forensically clean storage larger than the suspect system to be able to image the drive and conduct other valuable analysis.

Procedure: Retrieving Disk Information

Follow the procedure below to retrieve information about the disk drives located on a remote target.

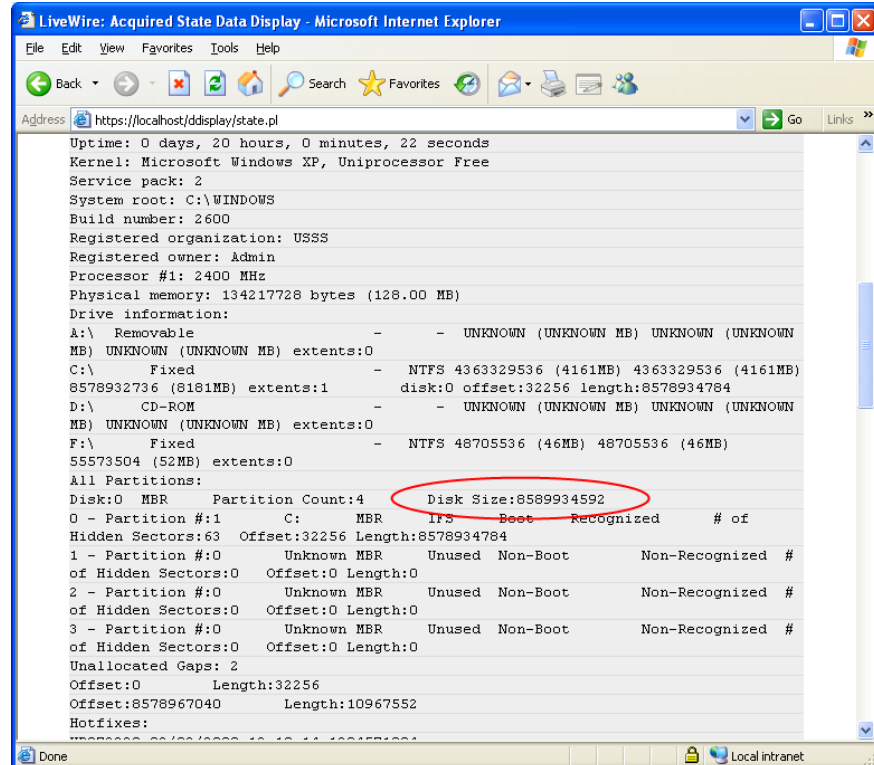
Step	Action
1	Click on the Data Display tab.
2	Beside Data to Display, use the pull-down menu to select General system information . Then click Display Selected State Info .



File System Status, continued

Procedure: Retrieving Disk Information, continued

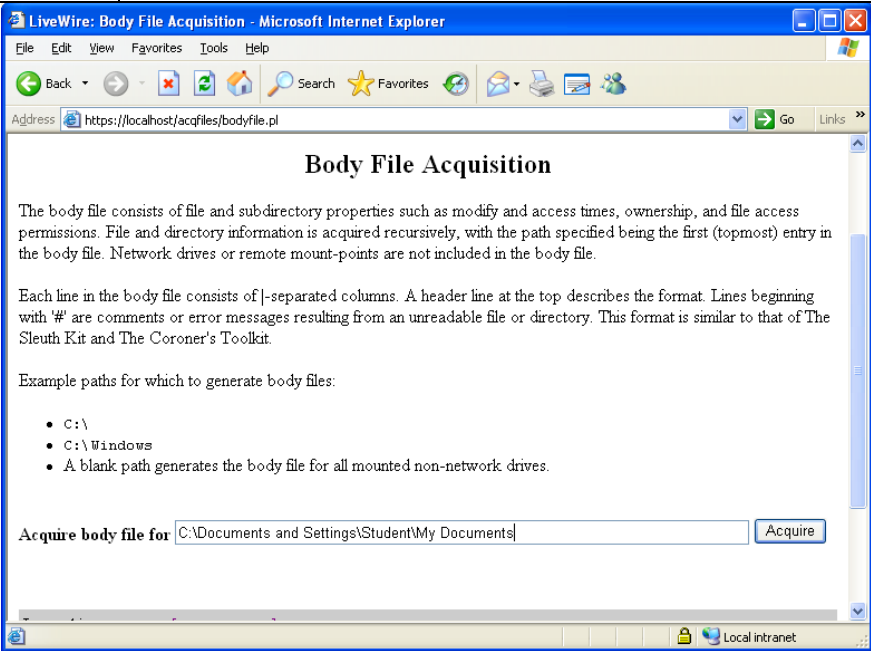
Step	Action
3	<p>Scroll down to view disk information. You can see that Disk0 is 8589934952 bytes in size. To convert that into a recognizable formation, you can use the following equation:</p> $8589934952 \text{ bytes} / 1024 / 1024 = 8192 \text{ MB}$ <p>Therefore, disk0 on 10.15.4.210 is 8 GB. Also, notice the logical disk information and RAM size are recorded. This information can be used to help determine the next course of action. Because imaging the entire drive will take a considerable more amount of time, the investigator may choose to only retrieve the logical partition, single files, etc. depending on the details of the investigation.</p>



File System Status, continued

Procedure: Body File Acquisition

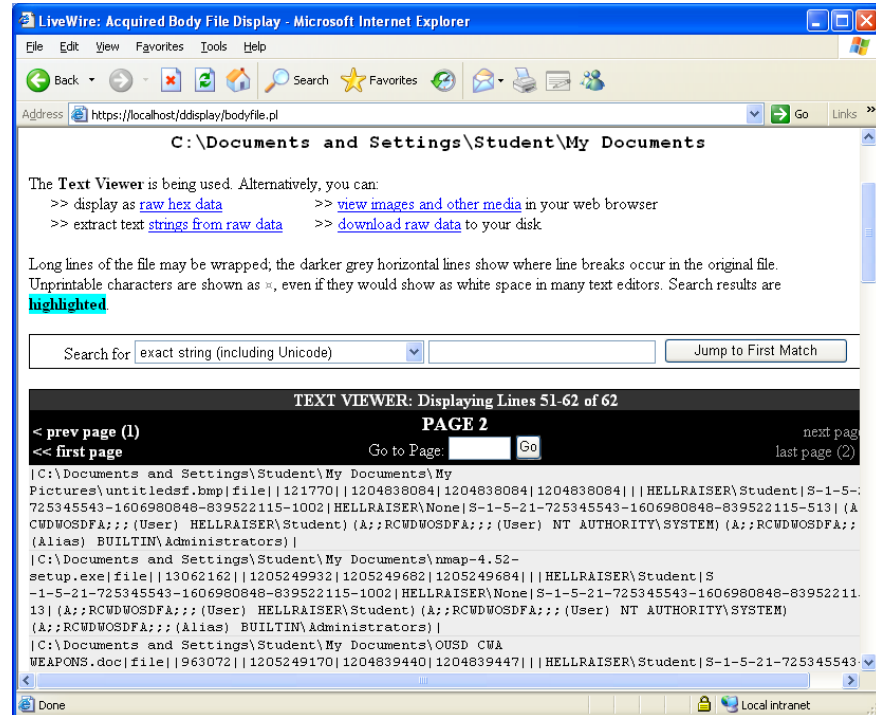
This section explains how to use LiveWire to retrieve information about files stored on the remote system. The body file acquisition gathers data about the files stored on the system, such as file modify and access times, ownership, permissions, etc. The output in this section is not easy to read, but could prove to be very valuable in the investigation and for archive purposes.

Step	Action
1	Click the Acquire Disk Data tab at the top of LiveWire Investigator.
2	Click Acquire a body file .
3	<p>Input C:\Documents and Settings\Student\My Documents in the input area.</p> <p><i>Even though there are spaces in the path, no quotes are necessary.</i></p>
	
4	Once the acquisition is successfully completed, click view the body file .

File System Status, continued

Procedure: Body File Acquisition, continued

Step	Action
5	A screen similar to the one below should be displayed.



Physical vs. Logical

Introduction

A physical image, which is used in forensic analysis, is a bit-for-bit duplicate of the hard drive in a system. In the event that a physical image cannot be obtained, the investigator collects a logical image, which only contains data from the active file system.

Physical Images

From a forensics standpoint, a physical image is preferable to a logical image because it may contain more evidence. Physical images:

- Contain information from the *entire* physical device or designated portion of it
- Are not file system-specific
- Capture all sectors within a designated area of a device, both in the system and data areas (including all files, unallocated space, swap space, etc.)
- Are typically placed in an Image file (a logical file that contains the bit-for-bit copy)

Examples of Physical Image Utilities

Both FTK Imager and the dcfldd utility, which are found on the Helix response CD, create physical images. These tools enable you to recover for analysis any deleted data or information that resided in slack space on the original drive. When a file is deleted and a smaller file is allocated to the same space, file slack may result. File slack can contain information from whatever previously occupied that space.

Physical vs. Logical, continued

Logical Images

In some cases, it may be necessary to collect a logical image instead of a physical one. You should be aware of limitations of logical images:

- Only contain information from the active file system
- Only contain enough information to reproduce logical volumes or parts of them
- Are file system-specific
- Allow registry and other system files to be backed up, but only if specifically requested
- Do not capture slack space, free space, or partition information
- Do not capture files that are open at the time of the image
- Do not capture any files that you do not have access to read
- Do not capture temporary files, such as pagefile.sys, win386.swp, etc.

Example of Logical Image Utilities

Microsoft's Windows Backup creates image files of an entire active file system. Because a logical image only contains active files, it is not possible to recover and analyze deleted files or slack space.

On-site Imaging Guidelines

When imaging on a crime scene, the same general principles for collecting evidence apply. However, you may face considerable time and material constraints. To add to that, you must also perform these actions correctly in a strange environment with unknown equipment.

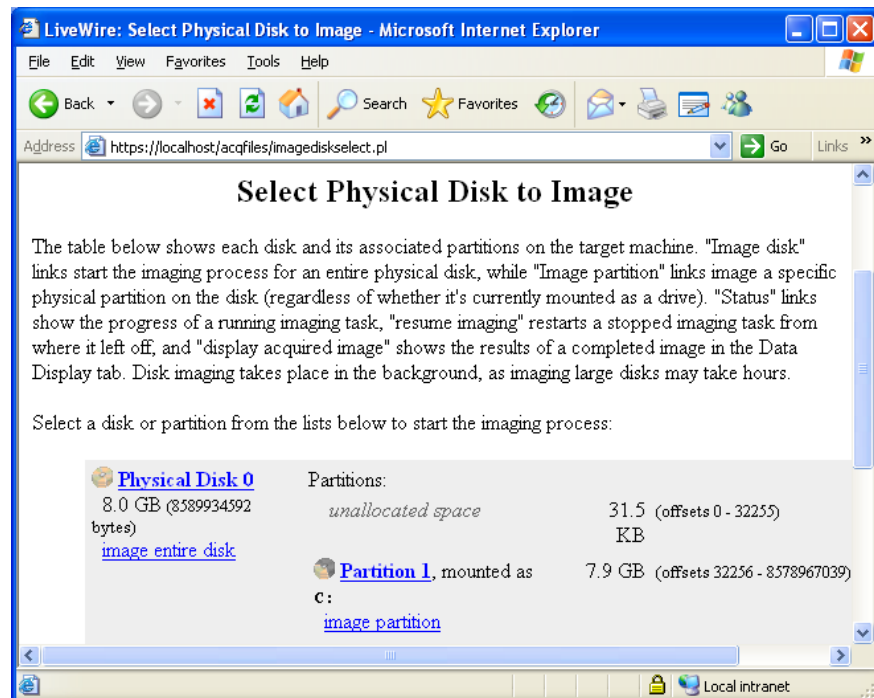
There will be challenges to accomplishing even simple actions, such as finding the appropriate settings in the BIOS or gaining access to the inside of the machine. Good preparation will help mitigate some of these problems. If you have to deviate from the general procedures for any reason, document the reasons why in your notes so that you can later explain in court.

Physical vs. Logical, continued

Procedure: Physical Disk Imaging

LiveWire has the ability to image an entire physical disk or an individual partition on that disk. This section will explain how to collect a physical disk image with LiveWire.

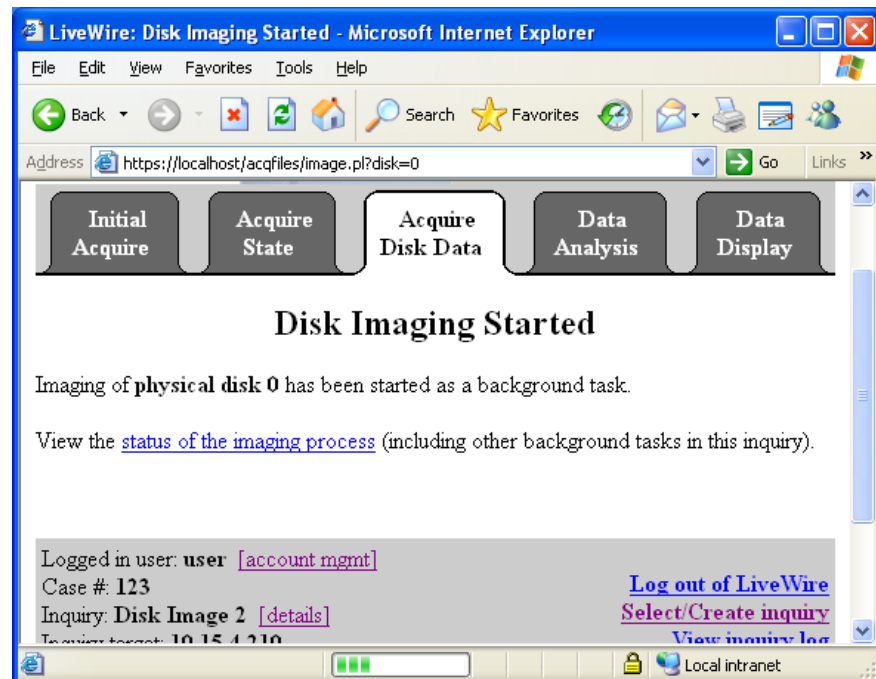
Step	Action
1	Click the Acquire Disk Data tab at the top of LiveWire.
2	Under the Raw Partition Data section, click Image a Physical Disk or Partition .
3	Click the Physical Disk0 link to begin imaging the disk.



Physical vs. Logical, continued

Procedure: Physical Disk Imaging, continued

Step	Action
4	The imaging process could potentially take a very long time. Therefore, LiveWire runs this process in the background to allow the investigator to carry out other tasks. To view the status, click the status of the imaging process link.



Physical vs. Logical, continued

Procedure: Physical Disk Imaging, continued

Step	Action
5	A new window opens which displays the progress of the imaging process with the option to end the job. This page will automatically refresh and can be closed.



Physical vs. Logical, continued

Procedure: Physical Disk Imaging, continued

Step	Action
6	To view the disk image, click on the Data Display tab.
7	Scroll down to the “Raw Disk & Memory images” section. In the pull-down menu, select Physical disk 0 and click Display Selected Disk Image .
8	The disk image will be displayed in the Hex Viewer with the same search capabilities.

LiveWire: Acquired Disk Image Display - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Address <https://localhost/ddisplay/image.pl> Go Links >>

Acquired Disk Image Display

Physical disk 0

The **Hex Viewer** is being used. This viewer was selected based on the data's contents. Alternatively, you can:

- >> display as [text](#)
- >> [view images and other media](#) in your web browser
- >> extract text [strings from raw data](#)
- >> [download raw data](#) to your disk

Each line shows 16 bytes of data beginning at the hexadecimal offset in the first column. The 16 bytes are in both hexadecimal and ASCII; unprintable characters are shown as x in the ASCII column. To search for raw hex values, use the regular expression search and precede each byte by \x; for example, to search for the pattern a0 00 09, use a regular expression of \xa0\x00\x09. Search results are **highlighted**.

Search for

HEX VIEWER: Displaying Bytes 0-799 of 8589934592

< prev page **PAGE 1** next page (2) >
 << first page Go to Page: last page (10737419) >>

0x0:	33 c0 8e d0 bc 00 7c fb 50 07 50 1f fc be 1b 7c	3 x x x x x P x x x x
0x10:	bf 1b 06 50 57 b9 e5 01 f3 a4 cb bd be 07 b1 04	x x x P x x x x x x x x x x
0x20:	38 6e 00 7c 09 75 13 83 c5 10 e2 f4 cd 18 8b f5	8 n x x u x x x x x x x x
0x30:	83 c6 10 49 74 19 38 2c 74 f6 a0 b5 07 b4 07 8b	x x x I t x 8 , t x x x x x x
0x40:	f0 ac 3c 00 74 fc bb 07 00 b4 0e cd 10 eb f2 88	x x x t x x x x x x x x x x
0x50:	4e 10 e8 46 00 73 2a fe 46 10 80 7e 04 0b 74 0b	N x x F x s * x F x x x x t x
0x60:	80 7e 04 0c 74 05 a0 b6 07 75 d2 80 46 02 06 83	x x x t x x x x u x x F x x
0x70:	46 08 06 83 56 0a 00 e8 21 00 73 05 a0 b6 07 eb	F x x V x x x ! x s x x x x
0x80:	bc 81 3e fe 7d 55 aa 74 0b 80 7e 10 00 74 c8 a0	x x x > x } U x t x x x x t x

Local intranet

Collection and Preservation

Overview

Every investigator knows the value of potential evidence and should be familiar with ways to preserve potential evidence at the scene. This section will explain how to use LiveWire to collect files from the remote system.

Preservation

You should alter the system as little as possible during your investigation by following sound first response principles.

Once on site, you should identify other devices that may have witnessed or captured information related to your investigation. Obtain any sniffer, router, firewall, and IDS (or similar) logs that may have captured traffic to or from the victim machine around the time of the incident.

Procedure: Collecting Files

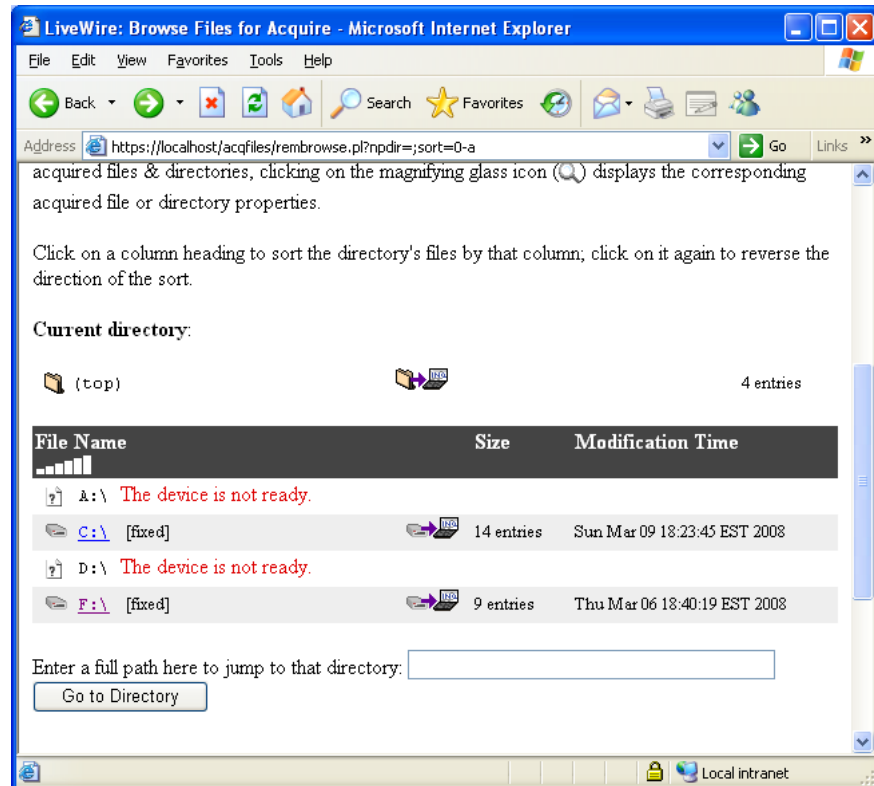
Follow these steps to use LiveWire to collect files from the remote system. This section continues from the previous section.

Step	Action
1	Click on the Acquire Disk Data tab at the top.
2	Click Browse and acquire files .

Collection and Preservation, continued

Procedure: Collecting Files, continued

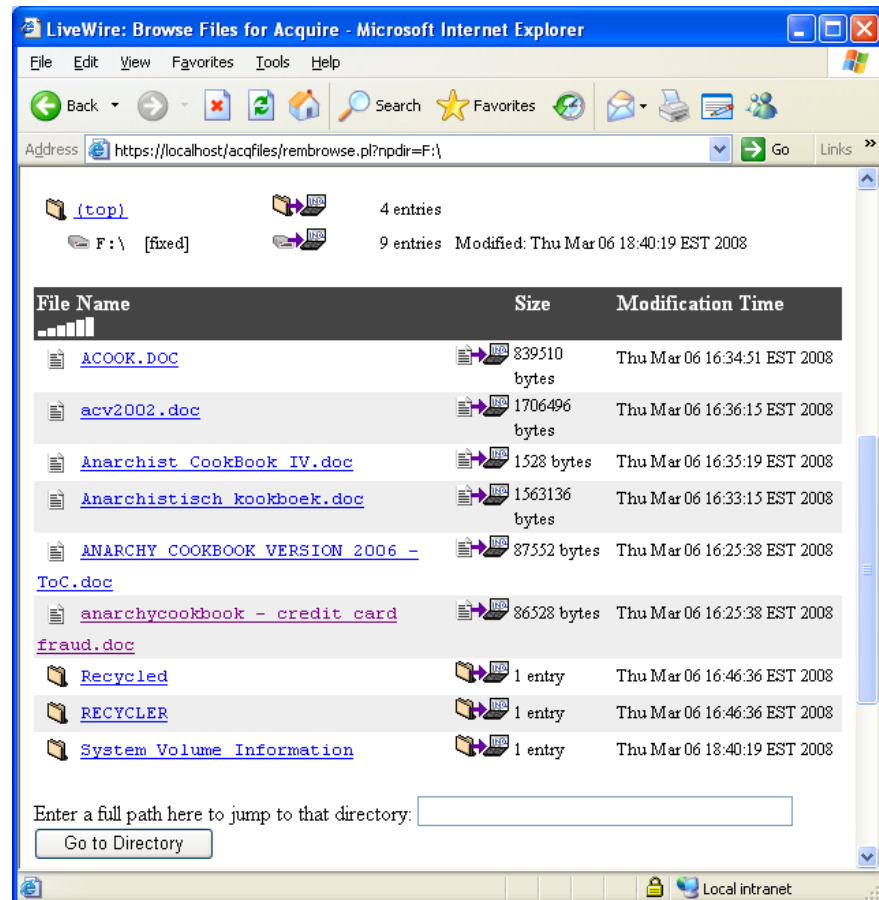
Step	Action
3	<p>This screen allows you to browse through the directory structure. Click the F:\ link.</p> <p><i>If you know the full path of the directory you want to view, the path can be entered in the area below the directory list.</i></p>



Collection and Preservation, continued

Procedure: Collecting Files, continued

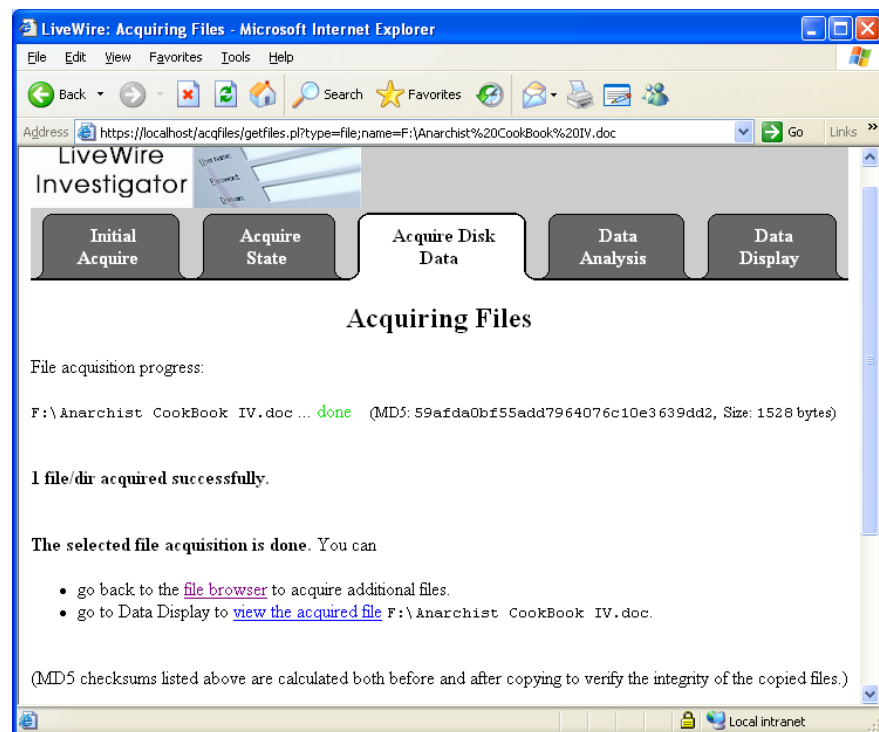
Step	Action
4	This view will show you the contents of the folder. Click the file Anarchist CookBook IV.doc to acquire the file.



Collection and Preservation, continued

Procedure: Collecting Files, continued

Step	Action
5	<p>The next screen will show the acquisition progress and automatically hash the file before and after this process. Two options are provided from this screen.</p> <ol style="list-style-type: none"> 1. Go back and acquire another file: Click file browser to return to the previous screen. 2. View the file: Click view the acquired file.



Hashing

LiveWire Hashing LiveWire uses MD5 (Message Digest 5) for hashing. MD5 creates a 128-bit message digest that is “unique” to the message. MD5 is currently the accepted standard for verification by the majority of the computer forensic community.

What is a Hash? A hash (or message digest) is a numerical value generated by applying a mathematical algorithm against a data set. Hashing algorithms will take variable length input and always output a “unique” fixed-length result. Essentially, it is analogous to fingerprinting an individual file.

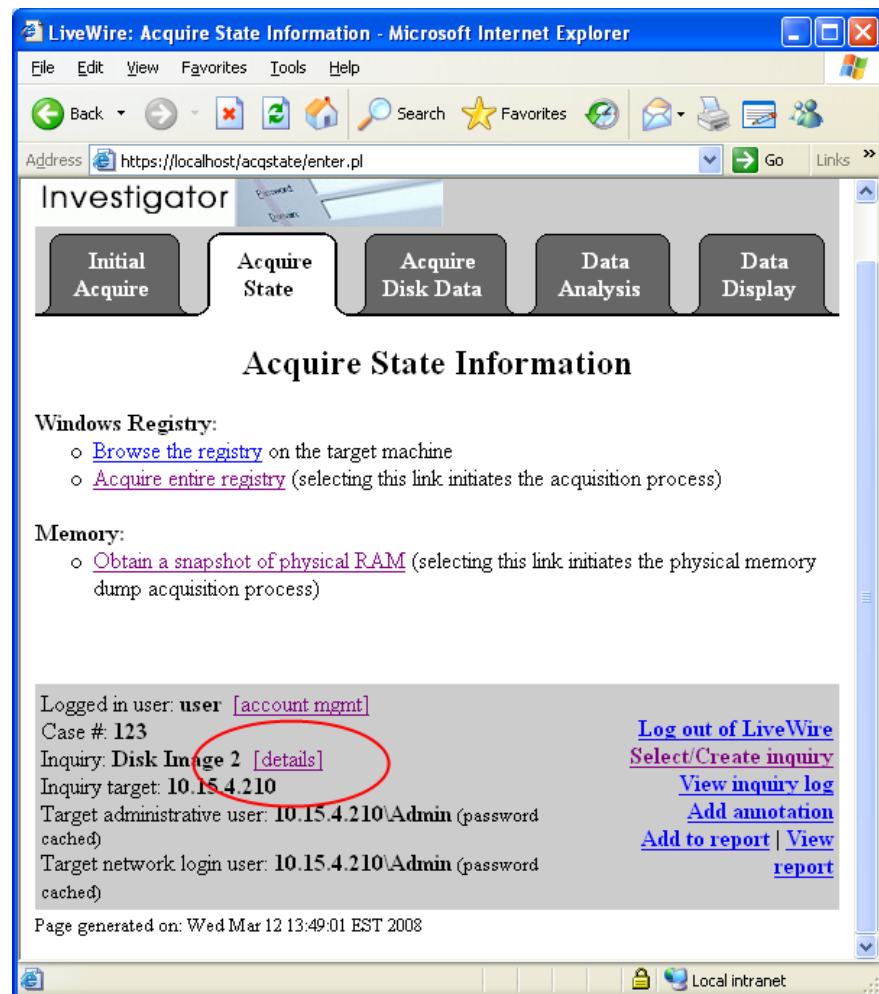
Once created, a hash value can be used to identify a file no matter where the file is found. As long as the file’s data does not change in any way, the same algorithm can be applied an infinite number of times and the resulting alphanumeric values will never change. If the hash value does change, it can be assumed that the file has been modified. Comparing hash values is an excellent way to check the integrity of files.

Hashing algorithms are “one-way.” This means that a hash can be created from file or device data, but you cannot recreate the data from the hash. Most importantly, it is nearly impossible to find two different data sets that naturally have the same hash value.

Hashing, continued

Procedure: Follow these steps to generate a hash for the inquiry.
Generating Hashes for the Inquiry

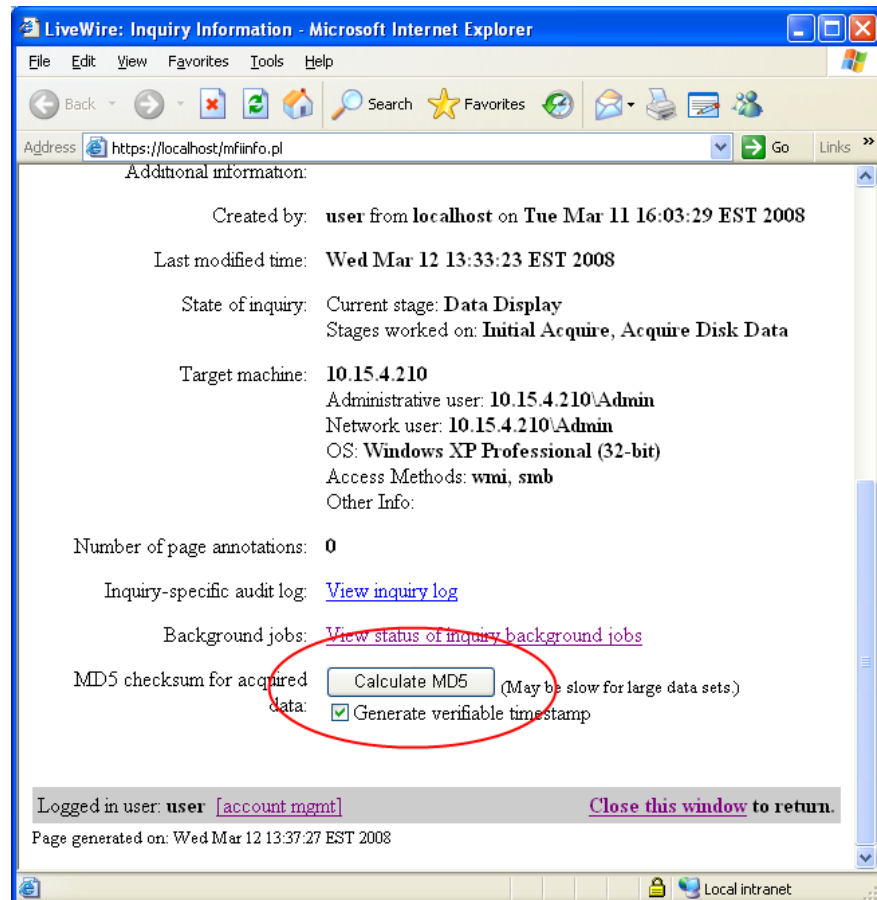
Step	Action
1	At the bottom of the page beside the Inquiry, click the details link.



Hashing, continued

Procedure: Generating Hashes for the Inquiry, continued

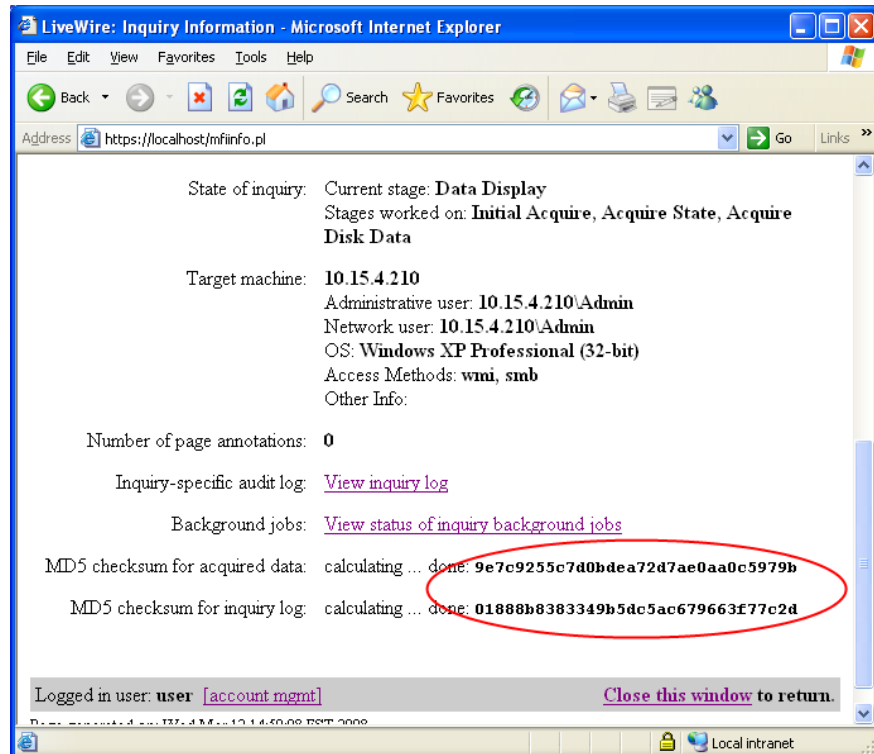
Step	Action
2	Scroll to the bottom of the Inquiry Information page. Click Calculate MD5 to create hashes for all acquired data.



Hashing, continued

Procedure: Generating Hashes for the Inquiry, continued

Step	Action
3	The size of the dataset will have a large impact on the amount of time it takes to complete the calculation. Once the hashing has completed, the hashes for both the acquired data and inquiry log will be displayed.



Hashing, continued

Procedure:
Viewing Acquired
File Hashes

Follow these steps to view the hashes for the files that have been acquired. These hash values can be compared to the hash values of the acquired files in the evidence drive to confirm the evidence is forensically sound.

Step	Action
1	Click the Data Display tab.
2	<p>Click Acquired Files and Directories towards the bottom of the page.</p> <p>This will display all the files that have been acquired from the target system along with their MD5 hash values.</p>

Lesson 6 – Malicious Code Analysis

Introduction

Windows systems are extremely vulnerable to attacks from malicious software known as malware. While some programs may not have a direct negative impact on the system, the applications present on a system could provide valuable information about the role and functionality of a suspect computer.

Purpose of this Lesson

The purpose of this lesson is to introduce the malware discovery function that is built into LiveWire.

Objectives

After successfully completing this lesson, you will be able to:

- Explain the malware search functions in LiveWire
- Conduct a malware analysis of a remote system

In this Lesson

The following table shows the contents of this lesson.

Topic	See Page
Malicious Program Search	12-94

Malicious Program Search

Overview

This lesson explains how to search for applications that could be categorized as malicious.

Types of Malicious Programs

There are many different types of malicious programs that may fall into any number of categories. Many of these programs may not necessarily show that the suspect machine was compromised but it could hint about the interests or the intentions of the user. Finding out the types of programs a user has on the machine could uncover what type of user owns the machine. Certain types of applications could tell the investigator how advanced the user may be which could guide the investigators search. For instance, if the person had an encryption program, such as TrueCrypt, that should raise the concern that there are encrypted volumes that may be hiding critical evidence.

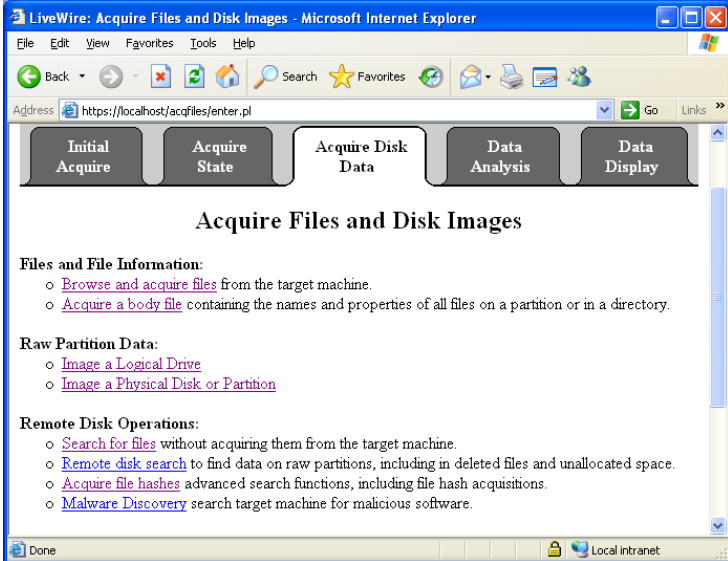
LiveWire performs malicious software scans by comparing hash signatures of files on the system against the National Software Reference Library (NSRL) database. The NSRL is a free database released by the National Institute of Standards and Technology (NIST). Using this information, LiveWire has the ability to search many different categories of malicious code. Some of these are:

- Anti Forensics
- Encryption
- Key Loggers
- P2P Tools
- Password crackers
- Rootkits
- Steganography
- Wireless

Malicious Program Search, continued

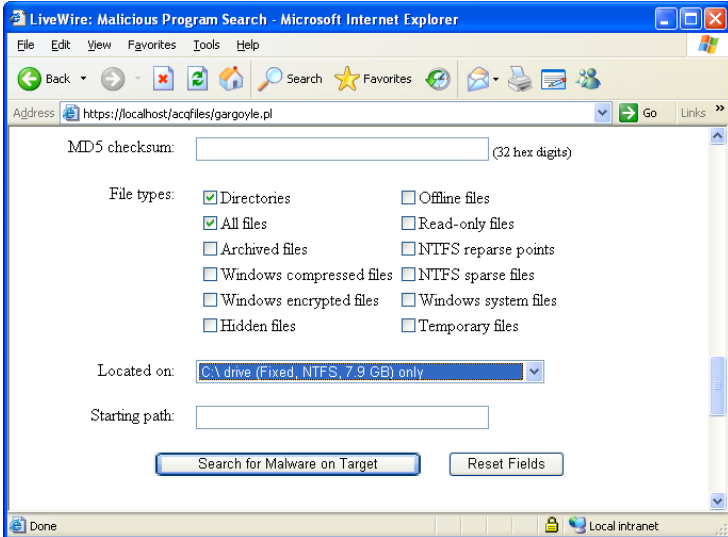
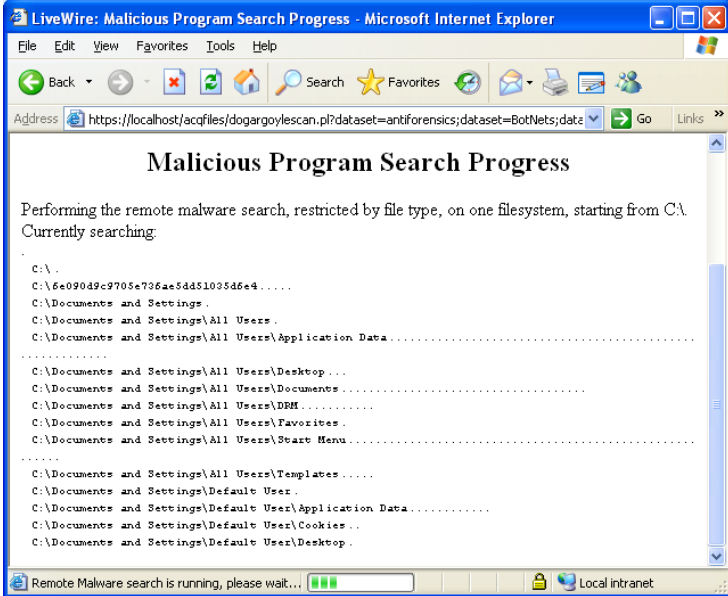
Procedure: Conducting a Malicious Program Search

Follow these steps to conduct a malicious program search of the remote machine. This lesson continues from previous lessons.

Step	Action
1	<p>Open the correct inquiry and click on the Acquire Disk Data tab.</p> <p>Note: Ensure you click the correct tab and do not get this confused with the very similar option under the Data Analysis tab.</p>
2	<p>Click on Malware Discovery towards the bottom of the page.</p> 

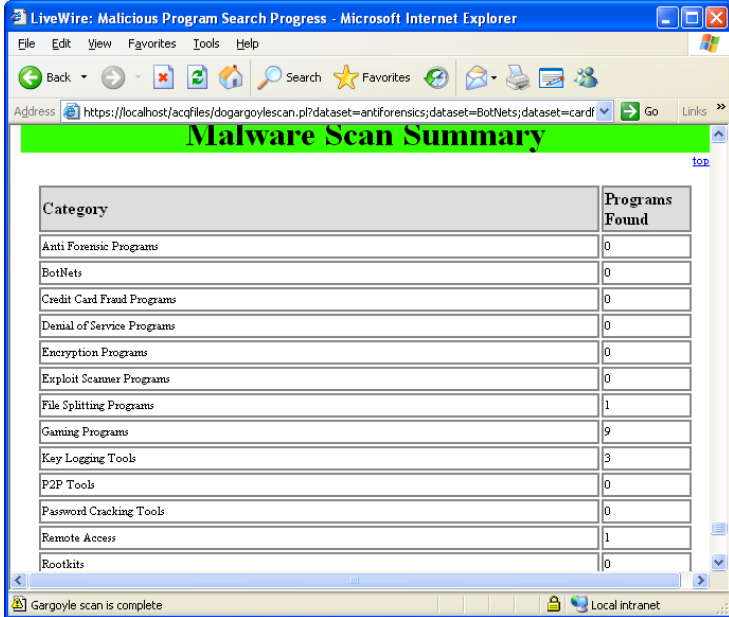
Malicious Program Search, continued

Procedure: Conducting a Malicious Program Search, continued

Step	Action
3	<p>Scroll down to the “located on” section. Select the C: drive only option. Notice the other options that are available. Then click Search for Malware on Target.</p> 
4	<p>A page similar to the one below will be displayed.</p> 

Malicious Program Search, continued

Procedure: Conducting a Malicious Program Search, continued

Step	Action																												
5	<p>Once the scan has completed, the Malware Scan Summary will be displayed. Click the View Gargoyle Report at the bottom of the summary.</p>  <table border="1"> <thead> <tr> <th>Category</th><th>Programs Found</th></tr> </thead> <tbody> <tr><td>Anti Forensic Programs</td><td>0</td></tr> <tr><td>BotNets</td><td>0</td></tr> <tr><td>Credit Card Fraud Programs</td><td>0</td></tr> <tr><td>Denial of Service Programs</td><td>0</td></tr> <tr><td>Encryption Programs</td><td>0</td></tr> <tr><td>Exploit Scanner Programs</td><td>0</td></tr> <tr><td>File Splitting Programs</td><td>1</td></tr> <tr><td>Gaming Programs</td><td>9</td></tr> <tr><td>Key Logging Tools</td><td>3</td></tr> <tr><td>P2P Tools</td><td>0</td></tr> <tr><td>Password Cracking Tools</td><td>0</td></tr> <tr><td>Remote Access</td><td>1</td></tr> <tr><td>Rootkits</td><td>0</td></tr> </tbody> </table>	Category	Programs Found	Anti Forensic Programs	0	BotNets	0	Credit Card Fraud Programs	0	Denial of Service Programs	0	Encryption Programs	0	Exploit Scanner Programs	0	File Splitting Programs	1	Gaming Programs	9	Key Logging Tools	3	P2P Tools	0	Password Cracking Tools	0	Remote Access	1	Rootkits	0
Category	Programs Found																												
Anti Forensic Programs	0																												
BotNets	0																												
Credit Card Fraud Programs	0																												
Denial of Service Programs	0																												
Encryption Programs	0																												
Exploit Scanner Programs	0																												
File Splitting Programs	1																												
Gaming Programs	9																												
Key Logging Tools	3																												
P2P Tools	0																												
Password Cracking Tools	0																												
Remote Access	1																												
Rootkits	0																												
6	View the report to see what details it provides. Note these options can be de-selected on the previous page.																												
7	If you would like to run the scan again with a different set of options, click the Data Analysis tab and select Malware Discovery. This option is only functional after an initial scan has been completed using the above steps.																												

This page intentionally left blank.

Lesson 7 – Alternate Data Collection Tools

Introduction Investigators need to be aware of other software that is acceptable to use during their investigations. In this lesson, you will learn about alternate data collection tools.

Purpose of this Lesson This lesson introduces alternative tools that could be useful to gathering information during an investigation.

Objectives After successfully completing this lesson, you will be able to:

- Discuss the functions of alternate tools
- Explain the functions of the Helix Live CD

In this Lesson The following table shows the contents of this lesson.

Topic	See Page
Windows Forensic Toolkit	12-100
Helix	12-103

Windows Forensic Toolkit

Overview

This section will introduce alternative tools that could prove to be a valuable addition to the investigator's toolkit.

Sysinternals - PsTools

The Sysinternals utilities, which were developed by Mark Russinovich, are capable of performing many different analysis functions on local or remote systems. The latest version of this collection can be downloaded from Microsoft.com.

psinfo.exe

Psinfo.exe will retrieve system information of the target system. Some of the data displayed is:

- Uptime
- Kernel version
- Product type
- Service pack
- Kernel build number
- Registered organization
- Register owner
- Install date
- IE version
- System root
- Processors
- Processor speed
- Processor type
- Physical memory
- Video driver

```
psinfo \\10.15.4.210 -u admin -p password
```

The output of all these commands can be redirected to a file by using the ">" string. To redirect the output of the psinfo command to a psinfo.txt file on the local machine:

```
psinfo \\10.15.4.210 -u admin -p password > psinfo.txt
```

Windows Forensic Toolkit, continued

pslist.exe

Pslist.exe will list the process currently running on the remote system.

```
pslist \\10.15.4.210 -u admin -p password
```

psloggedon.exe

Psloggedon.exe will display a list of currently logged on users on the remote system for both local and remote users.

```
psloggedon \\10.15.4.210
```

psexec.exe

Psexec.exe is an advanced utility used to execute commands on a remote system. It also has the ability to copy a program from the local system to the remote target and then execute that program interactively.

This command will connect to 10.15.4.210 using the username “admin” and the password “password” and then run the “cmd” command. This connects to the remote system with a terminal session which will allow you to execute all commands on the remote system from your local console.

```
psexec \\10.15.4.210 -u admin -p password cmd
```

psfile.exe

Psfile.exe is used to view files that are opened remotely on the target system. This command will not display files that are locally opened on the target system.

```
psfile \\10.15.4.210 -u admin -p password
```

psgetsid.exe

Psgetside.exe will retrieve the SID of the target system.

```
psgetsid \\10.15.4.210 -u admin -p password
```

Windows Forensic Toolkit, continued

psloglist.exe

Psloglist.exe will retrieve the logs from the target system. By default psloglist.exe will show the contents of the *system* event log. The application, security, or other log can be retrieved if specified.

```
psloglist \\10.15.4.210 -u admin -p password
```

```
psloglist \\10.15.4.210 -u admin -p password application
```

```
psloglist \\10.15.4.210 -u admin -p password security
```

psservice.exe

Psservice.exe will retrieve a list of running services on target system.

```
psservice \\10.15.4.210 -u admin -p password
```

Helix

Introduction to Helix

Helix, a customized distribution of Knoppix created and maintained by e-fense, Inc., is geared toward forensics and incident response. First released to the public in November 2003, Helix was first created to be used as an internal tool for incident response and forensics to create forensically sound images.

The customizations of Helix have been made to forensically prevent the CD from altering data on the host computer.

Helix has been created with two different operating modes. There is a Windows mode and a Linux mode.

The latest version of Helix can be downloaded at:
www.e-fense.com/helix/downloads.php

Helix - Windows Mode

The Helix Windows mode is created with Windows executables and contains many tools for incident response on a Windows machine. In this mode, the CD runs standard Windows applications to gather information from a ‘Live’ running system. This can be useful where systems cannot be shut down or where potential evidence would be destroyed by taking the system offline.

Note: When a system is up and running it is constantly changing. Running Helix in the live environment will make changes to the system. It is important to be aware of this known fact and that it is documented and understood. When Helix is first opened, the following warning message will be displayed. You must click accept to continue.

Helix, continued

Helix - Windows Mode, continued



The next screen is the first default screen that is displayed to the user. There are many tools located on the CD that can be used to gather volatile and non-volatile information.



Helix, continued

Helix - Windows Mode, continued

The Windows mode on the Helix Live CD has the ability to acquire images of a live system. This can be done using the Live Acquisition feature of the CD. It is possible to image the physical memory, physical drive, or logical partitions. The images can be saved to an attached device, network share or to an evidence capture machine using NetCat.



FTK Imager is also available on the CD for creating forensic disk images as well as saving them in different formats, such as raw dd images and E01 (EnCase) images. It is located on the menu bar under Quick launch > FTK Imager. FTK Imager also allows for imaging physical and logical drives.

List of Some Available Tools in Windows Mode		
Command Shell	FTK Imager	Sys Info Viewer
Drive Manager	Win Audit	Zero View
Per-Search	WFT	NetCat
VNC Server	PuttySSH	File Recovery
Rootkit Revealer	Screen Capture	Password Viewers

Helix, continued

Helix - Linux Live CD Mode

The Linux mode of Helix is a pure Live CD that allows for the “dead box” forensics. This allows the user to investigate a computer system without forensically changing any data on the hard drive. Many tasks can be carried out with Helix, such as the ability to forensically duplicate disks as well as analyze those forensic disk images.


To start Helix in Linux mode, the system will need to be booted to the bootable Helix CD. The BIOS on the motherboard must be correctly configured.

Helix – Linux Live CD Disk Mounting and Imaging

When Helix is booted into Linux mode, it will automatically mount all storage devices in read-only mode. It will also mount devices with the noatime option, which will prevent any change to the access times of files stored on the disk. Although Helix will mount drives as read-only by default, Helix can be forced to mount devices as read-write by typing:

```
mount -rw <device> <mount point>
```

Once Helix has fully booted, the screen will appear similar to the image below. It will show a list of all storage media mounted on the left side of the screen. The taskbar is located at the bottom of the screen and the “Start” menu is an icon on the Helix CD cover.



Helix, continued

Helix – Linux Live CD Disk Mounting and Imaging, continued

Several GUI utilities, such as Adepto, Air, and Linen, are included on the disk for creating forensically sound images.



These tools can be used to create and store those images in various locations, such as a network share, locally attached storage drives, and even across the network to an evidence collection machine. This provides many options for data collection. In order to ensure you do not overwrite your evidence media, always be aware of exactly what drives are being imaged and where they are being imaged to.

Helix, continued

Helix – Linux Live CD Forensic Tools

Helix also provides tools for investigating collected disk images. Autopsy, as shown in the image below, is a popular Linux tool for viewing and searching images. Autopsy is the GUI interface to a suite of command line forensic tools named The Sleuth Kit.

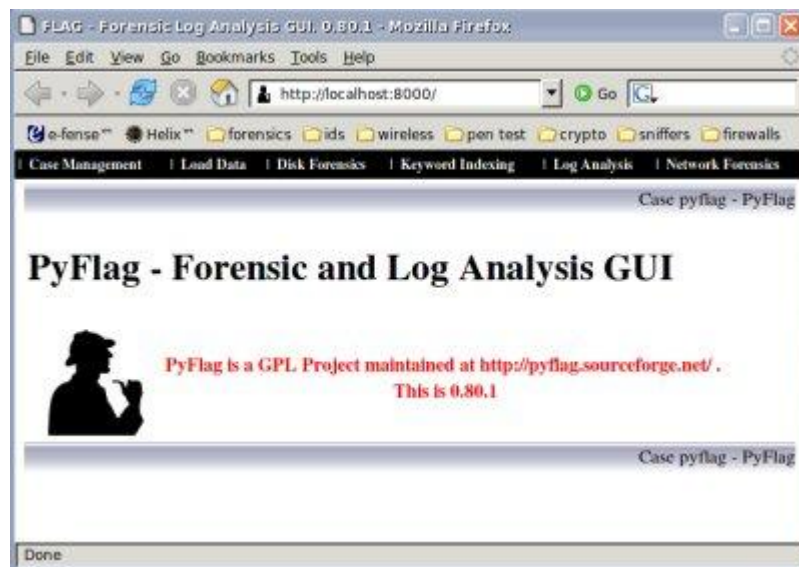


Helix, continued

Helix – Linux Live CD Forensic Tools, continued

PyFlag, a forensic and log analysis application created by the Australian Department of Defense, is also available on the Helix disk. PyFlag uses a backend database to assist managing large amounts of data. This tool is Web driven and can be deployed on a central server, which allows several users to use the tool at the same time.

PyFlag is able to examine forensic evidence from disk images, logs, and network captures.



Live CD Benefits

Live CDs can be very useful for testing, evaluating, or learning without the need of dedicated hardware. Many Linux distributions are available as Live CDs and can be freely downloaded from the Internet. Each is designed with a specific task in mind.

Helix is just one great example of a feature-rich Live CD created for forensics and incident response. Helix provides multiple options for investigations whether the suspect machine is turned off or is up and running.

This page intentionally left blank.

Appendix A – Intrusion Report Template

Introduction

Responding to a network incident is usually a very complex and intricate experience. For even the smallest of incidents, there is a large amount of information that needs to be collected and documented to successfully analyze the incident. The United States Secret Service has released a standardized network incident report template to expedite the collection of network intrusion information.

Purpose of this Appendix

The purpose of this appendix is to introduce you to the United States Secret Service's Network Intrusion Report template which can be used to document an ongoing network intrusion response. You will learn about the benefits of using this template as well as where to find it.

Objectives

After completing this appendix, you will be able to:

- Explain the usefulness and need for an intrusion report
- Download the States Secret Service's Network Intrusion Report template

In this Appendix

The following table shows the contents for this appendix.

Topic	See Page
USSS Electronic Crimes Network Incident Report	A-2

USSS Electronic Crimes Network Incident Report

Overview

Many inexperienced responders simply do not know the full extent of information to collect or to request from witnesses on the scene. To alleviate many of these problems, the United States Secret Service has released a standardized network incident report template.

This template can be found at the Forward Edge II Web site, <http://www.forwardedge2.com>. More specifically, the actual form can be downloaded at:

<http://www.forwardedge2.com/pdf/form-in.pdf>

Filling Out Report

The Network Incident Report is used to request assistance in a network incident from a local USSS Electronic Crimes Task Force (ECTF). By filling out the report in completion, a responder can frame the full extent of a network crime to ensure that the proper support is provided.

However, the use of this form does not obligate the responder to request for assistance. The standardized approach to the Network Incident Report allows for it to be used as a guideline for current and future incidents. It acts as an aid to the incident responder to ensure that all information is collected and filled out accordingly.

Appendix B – Volatile Data Collection

Introduction

This appendix introduces the collection of volatile data and focuses specifically on the collection of data from a Windows system. Not all forensic tools will work exactly the same on all versions of Windows. Service packs, updates and security features may impact how tools interact with the system. However, the forensic methodology will remain the same no matter what system you are examining.

Purpose of this Appendix

The purpose of this lesson is to provide investigators with the ability to retrieve volatile data before shutting down a live system.

Objectives

After completing this lesson, you will be able to:

- Explain the importance of the collection of volatile data
- Use the Helix disk to collect information

In this Appendix

The following table shows the contents of this appendix.

Topic	See Page
Overview of Volatile and Non-Volatile Data	B-2
Introduction to Helix Live CD	B-3
Collecting Volatile Information	B-5
Imaging Encrypted Volumes	B-12

Overview of Volatile and Non-Volatile Data

Introduction

Before pulling the plug and imaging physical drives, there are times when it might be beneficial to gather data from a live system. This data includes:

- Volatile data: Data that would be otherwise lost when the system is shut down
- Non-volatile data: Data such as the size and number of volumes in the system.

What is Volatile Data?

Volatile information is data that will be gone once power is removed from the system. Among other things data can include:

- Current network sessions
- Current ports and services open on the system
- Current processes running on the system

What is Non-Volatile Data?

There is some information that you might want to collect and view onsite to help determine the best way to image the system. This information is non-volatile and can include:

- Size of the victim hard drives, hence helping you decide if you need to image a disk drive per partition or the whole drive at once
- Number of disk drives on the victim machine to ensure that all are imaged

Introduction to Helix Live CD

Introduction to Helix

Helix is a customized distribution of Knoppix geared toward forensics and incident response. Created and maintained by e-fense, Inc., Helix was first created to be used as an internal tool for incident response and forensics to create forensically sound images. Helix was first released to the public in November 2003.

The customizations of Helix have been made to prevent the CD from altering data on the host computer.

Helix has been created with two different operating modes: a Windows mode and a Linux mode.

Helix - Windows Mode

The Helix Windows mode is created with Windows executables and contains many tools for incident response on a Windows machine. In this mode, the CD runs standard Windows applications to gather information from a “live” running system. This can be useful where systems cannot be shut down or where potential evidence would be destroyed by taking the system offline.

Note: When a system is up and running it is constantly changing. Running Helix in the live environment will make changes to the system. It is important to be aware of this fact and that it is documented and understood.

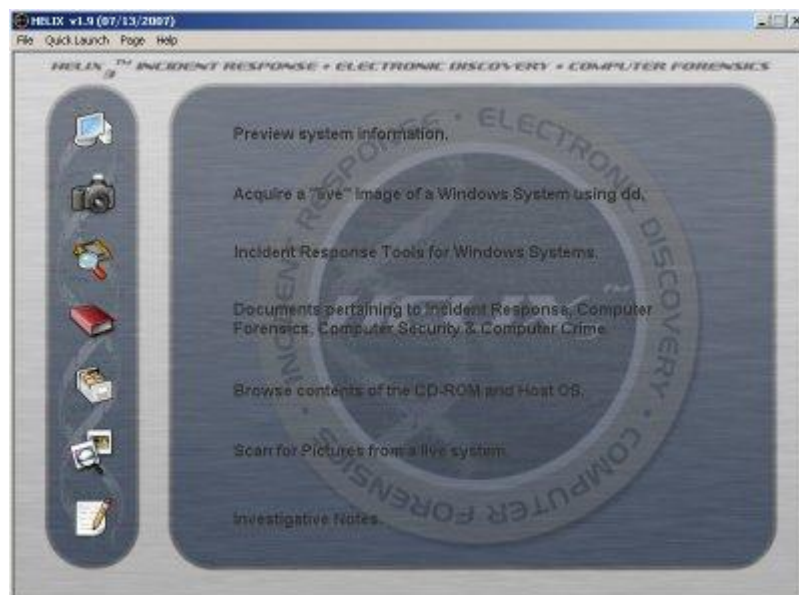
Introduction to the Helix Live CD, continued

Helix - Windows Mode, continued

To use Helix, place the disk in the target system. The first screen that will appear will be the Warning screen as depicted below. You must click accept to continue.



The next screen displays a number of options in the form of icons on the left side of the screen from which the user may choose. There are also options available on the toolbar.



Collecting Volatile Information

Imaging RAM

While many volatile data collection commands will retrieve obvious bits of information, they only grab small portions of what is available from within the computer's memory. They will not grab other evidentiary items that may be in memory, such as information found on open Web pages and running within open applications.

You should image the RAM of a computer early in the volatile data collection process. As additional commands and processes are run later, the data from these commands and processes can overwrite critical information contained in RAM. Imaging the RAM beforehand ensures that the information collected is the same as its original state.

Using a variety of methods, it is possible to image the complete contents of a system's memory for later analysis. There is no best way to imaging RAM; every method has issues. One of the ways to completely grab all memory is to initiate a system crash, which dumps all of the RAM's contents into a local dump file.

Obviously, as this method actually crashes the computer, it is not a preferred method for responders. Instead, you can use the dd utility to image the physical memory. The only negative aspect to this method is that RAM will be continually changing and updating during the imaging process.

Procedure: Imaging RAM


Following the steps below to image the system's physical memory using the dd command found within Helix.

Step	Action
1	From the Helix main menu, use the pull-down menus to select Quick Launch > Command Shell.
2	From the newly opened command line terminal, image the RAM by typing the following command line in one line, replacing "D:\\" with the drive and folder of your evidence repository: <code>dd if=\\.\PhysicalMemory of=D:\RAM.dd conv=noerror</code>
3	The process will take a number of minutes to completely image all of the memory. You will occasionally see error messages stating that physical memory ranges could not be read. These refer to memory ranges that are locked, and can be disregarded.

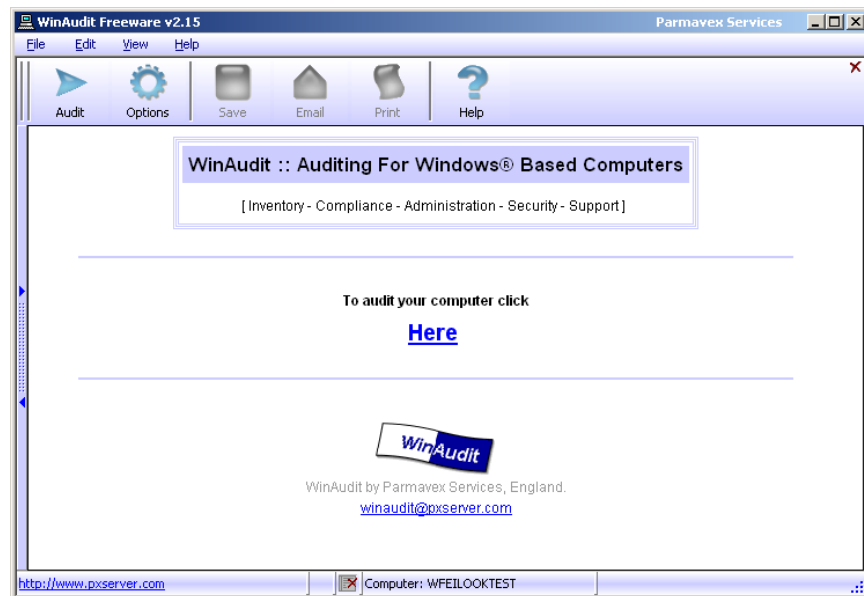
Collecting Volatile Information, continued

Helix - Windows Mode, continued

To acquire information on a Windows system in a quick and easy manner while limiting the impact on the target system, select the Quick Launch option on the toolbar. From the drop-down menu select Win Audit. The following window will display.



Click Yes to execute the program. The following screen will be displayed.

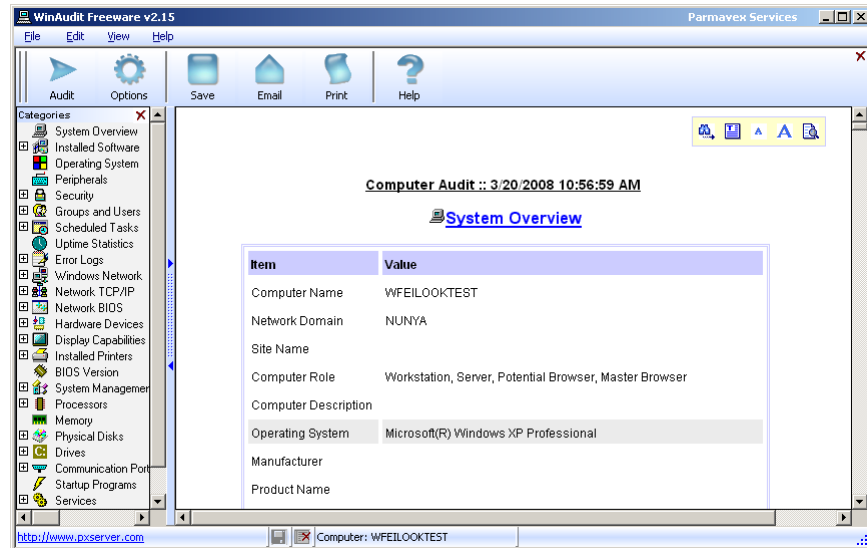


Click on the Here link to start acquiring data.

Collecting Volatile Information, continued

Helix - Windows Mode, continued

The program should take just a few moments to query the target system for information. Once it completes a screen similar to the following will be displayed.



Two clearly distinguishable panes will be evident. The left pane is labeled “Categories.” In that pane will be a number of options that when selected will display corresponding information in the right pane. For instance, the first option in the left pane is System Overview. By default, it is selected when the window is first displayed. As can be seen in the screenshot above, the corresponding information is available for scrutiny in the right pane.

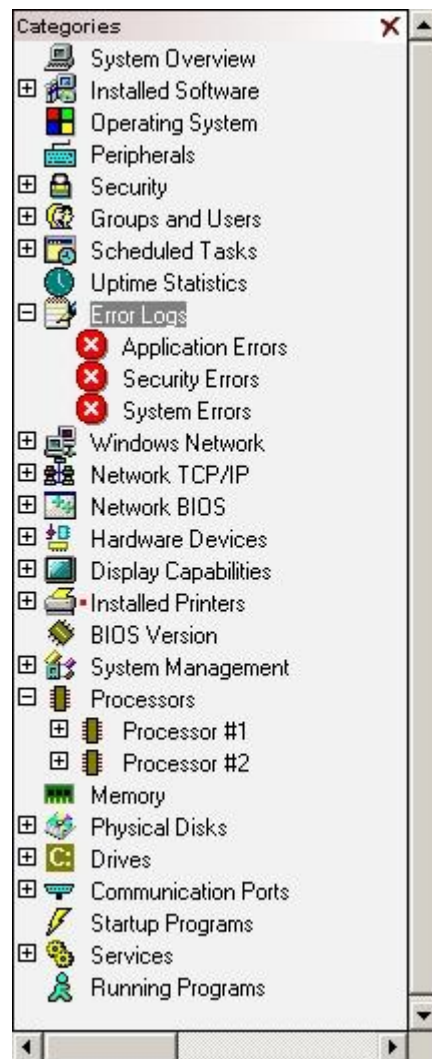
Information acquired during the audit will be placed in one of the categories depending on to what it pertains and from where it was obtained.

Collecting Volatile Information, continued

Helix - Windows Mode, continued

The following screenshot depicts the categories created from running Win Audit.

Many of the categories in the left pane can be expanded to provide access to additional information. As seen in the example below, the Error Logs option, when expanded, will display three additional options, one each for the System, Application and Security event logs.



Again, when a category is selected in the left pane the corresponding information will be displayed in the right pane.

Collecting Volatile Information, continued

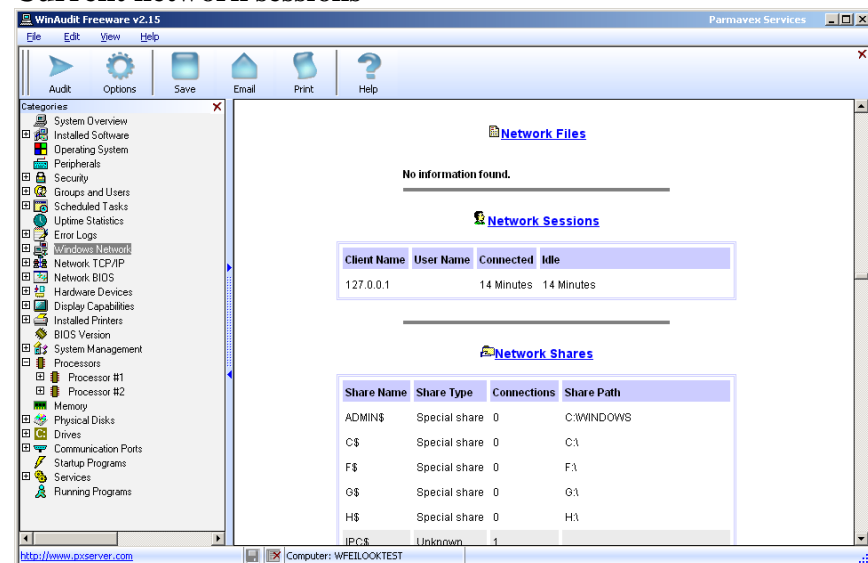
Helix - Windows Mode, continued

The application acquires some of the most commonly sought information during an initial or first response. Among other information, this includes:

- Current network sessions
- Open files
- Open ports
- Active processes
- Running programs

The following screenshots illustrate some of these categories as displayed in Helix.

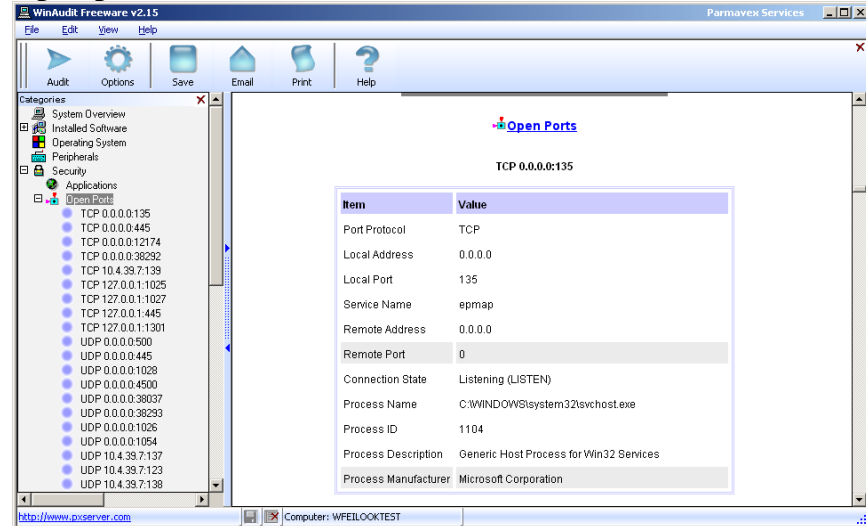
Current network sessions



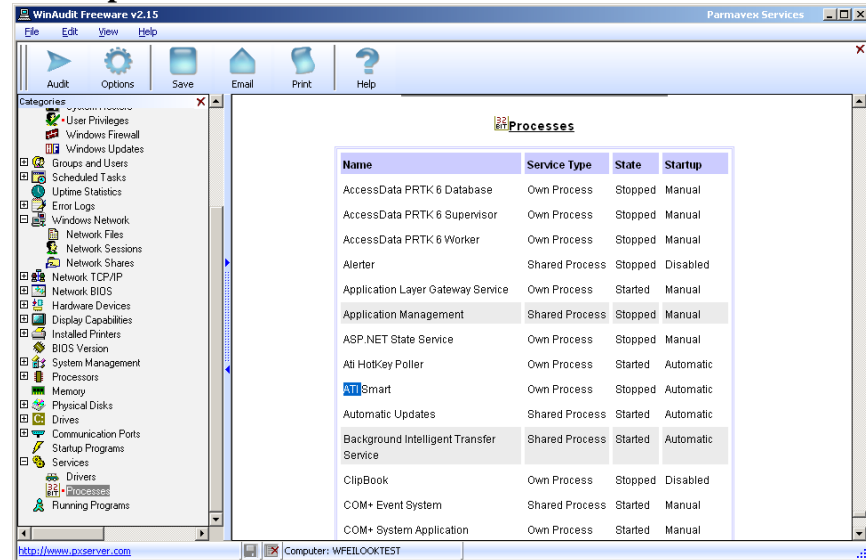
Collecting Volatile Information, continued

Helix - Windows Mode, continued

Open ports



Active processes



Collecting Volatile Information, continued

Helix - Windows Mode, continued

Once acquired, the information can be saved to whatever location is desired for the storage of evidential items. In order to do so, click on the File option on the toolbar, go to a storage directory and save the output.

By default, Helix saves the output in HTML format. It creates three HTML files. One of the files contains information from both panes of the main Win Audit screen. The other two HTML files contain only information relating to either the categories (left) pane, or the information displayed in the main display (right) pane of the Win Audit window. This affords the examiner three different options for analyzing the output. The examiner can select whichever option is most convenient or appropriate depending on requirements.

Output can be saved in other formats, including PDF, text and as a database file. Additionally, there is also an option available through the File option on the toolbar for sending the output to another location via e-mail. This could be extremely advantageous. For instance, suppose the situation is one in which an intrusion is suspected but it has not yet been confirmed. An initial responder who is not experienced with intrusion detection could respond and use Helix to collect data. It could then be e-mailed to an experienced examiner who could make a decision as to whether or not it would be beneficial to image the entire system. This could save time and resources.

Helix provides automated tools that require little experience or expertise to run. As you become more familiar with forensic methodologies, you may choose to create your own forensic tool kits and use less intrusive methods to acquire volatile data. The primary goal is to acquire the desired data in the least intrusive manner.

Imaging Encrypted Volumes

Overview

In recent years, the frequency of encrypted volumes has grown in both the consumer and corporate environments. Numerous user-friendly applications are available to create an encrypted volume for mounting at any time. One such program is TrueCrypt. TrueCrypt allows users to create a large, encrypted image file which can be mounted to store files.

When in use, the encrypted file is mounted as a new drive letter accessible by all users. A password must first be entered to open the image file, but afterwards it is completely open to the entire system for access. When the encrypted file is not in use, then all of the data remains secure within the encrypted volume.

When responding to a live machine and performing volatile data analysis, it is important to determine if a volume encryption application is active and running on the system. If so, care should be taken to identify if any encrypted volumes are open on the system.

If encrypted volumes are open on the system, the logical volume should be imaged to your evidence repository. Failure to do so will prevent an investigator from ever being able to access the file without the volume's password.

Procedure: Imaging a Logical Volume

Follow the steps below to image a logical volume of the local system to your evidence repository.

Step	Action
1	From the Helix main menu, use the pull-down menu to select Quick Launch > Command Shell.
2	From the newly opened command line terminal, image the RAM by typing the following command line, replacing "G:" with the drive letter of the logical volume and "D:\" with the drive and folder of your evidence repository: <code>dd if=\\.\G: of=D:\VolumeG.dd conv=noerror</code>

Appendix C

Understanding Computer Hardware

- Overview** This module explains the procedures necessary for safe handling of computers. Students will learn the primary hardware components that power the data processing and storage functions of every computer. An understanding of motherboards, CPUs, memory, and bus is essential to knowing how a computer system works.
- Purpose of this Module** The purpose of this module is to provide students with an understanding of primary computer hardware components.
- Objectives** After successfully completing this module, you will be able to:
- Practice safety procedures when handling computer equipment
 - Identify major computer components
 - Identify and explain motherboard types
 - Recognize individual motherboard components including chipsets, jumpers and switches, power supply and connections
 - Define Basic Input/Output System (BIOS)
 - Recall CPU functions and memory
- In this Module** The following table shows the contents of this module.

Topic	See Page
Lesson 1 – Safety Briefing	C-3
Lesson 2 – Overview of Computers	C-5
Lesson 3 – Motherboards and Components	C-17
Lesson 4 – CPU and Memory	C-39

This page intentionally left blank.

Lesson 1 – Safety Briefing

Introduction To ensure the well being of each student, a safety briefing is given before students begin using the classroom computers.

Purpose of this Lesson You will learn the procedures necessary for safe handling of computers.

Objectives After successfully completing this lesson, you will be able to:

- Identify the steps to take to protect yourself from injury when using a computer
- Explain how to protect computer components and stored data

In this Lesson The following table shows the contents of this lesson:

Topic	See Page
Safety Briefing	C-4

Safety Briefing

The Need for Safety Procedures

During the NITRO course, you will perform several practical exercises involving the disassembly and reassembly of computer components. Therefore, it is imperative that you follow the safety procedures presented here. You will be given a wrist grounding strap and electrostatic mat to use in the classroom.

Warning All electrical devices contain components that may injure or kill people who do not take proper safety precautions.

Step-by-Step Safety Procedures

Step 1: Turn off power and disconnect main power cables.

Before opening a computer or handling any component, always ensure that all computer devices are turned off and the main power cables are disconnected.

Warning To avoid death or serious injury, never open a power supply or monitor chassis. Capacitors inside a monitor hold electrical charges even when the monitor is unplugged. Therefore, it is important not to open a monitor chassis. If a problem exists with either device, take it to a professional.

Step 2: Use a wrist grounding strap and electrostatic mat.

Static electricity can damage or destroy most computer circuitry, cards, and memory. To protect against static discharge, use a wrist grounding strap and an electrostatic mat when handling computer equipment. If a strap is unavailable, first touch a metal object prior to handling computer components.

Step 3: Remove all jewelry from hands.

Before handling the computer, remove all jewelry from your hands including watches, rings and bracelets. Objects like watches, rings, and bracelets are good electrical conductors. Jewelry may also get caught in computer components and ruin them.

Lesson 2 – Overview of Computers

Introduction

This lesson presents the key components of the computer, history of computing, and basic terminology. Computer components are identified and their roles are reviewed in relation to the computer system as a whole.

Purpose of this Lesson

You should be familiar with the history of computing, know the basic computer components, and understand the terminology used to describe system components.

Objectives

After successfully completing this lesson, you will be able to:

- Identify the basic computer components
- Define basic computer terminology
- Explain the history of the modern computer

In this Lesson

The following table shows the contents of this lesson.

Topic	See Page
Introduction	C-6
History of Computers	C-9
Basic System Components	C-12

Introduction

What is a Computer?

A computer is a machine that performs high-speed operations and processes data. In its simplest form, a computer is a large collection of electronic switches, or transistors, operating very quickly in a specific order. Programs tell transistors how, when, and in what order to turn on and off. These on and off actions are equated to the binary system of numbers 1 (On) and 0 (Off). These 1s and 0s are stored in computer systems as bits and make up the building blocks of all data and information. A set of 8 bits is used to create a single byte of information, such as a character or number.

Computer programs consist of streams of bits, each bit indicating on or off. These streams are called a data stream or a bit stream. Computers can only recognize information in bits called machine language. Software and hardware translate these numerical streams into a human readable format.

Computers range in function from the general purpose desktop PC to massive mainframes to specialized chips in children's toys. The most prevalent is the PC, which usually consists of a case that sits on the floor or desk, a monitor, a keyboard, and various peripherals like printers.

PC functions range from general purpose, stand-alone systems to specialized servers that perform networking functions.

Introduction, continued

Laptop and Notebook Computers

Laptop and notebook computers are designed to be portable. Early models were heavy, slow, and did not have the same storage capacity as their desktop counterparts. Today, these systems rival the performance of most desktop PCs.

- **Laptops:** Laptops typically weigh seven pounds or less and are approximately 9x12x2 inches in size. They are powered by rechargeable batteries and AC adapters. Laptops can offer high performance and multimedia capabilities. A docking station can be added to enable connectivity to networks, regular monitors, keyboards, and other peripherals.
- **Notebooks:** Notebooks are smaller and lighter than a laptop. In general, they lack the high-end multimedia functions of the laptop. Yet, many notebooks have comparable hard drive and memory configurations and are equipped with sound and CD-ROM drives.

Personal Digital Assistants (PDAs)

PDAs, also known as palm pilots, IPAQs, and pocket PCs meet the demand for a reduced-function portable computer. PDAs enable users to manage files and to swap data with a desktop computer. Most are used to maintain contact lists and to track appointments. Current models can help manage e-mail, paging, and faxes. Some have wireless connectivity to other devices using infrared connections. Others can connect to the Internet through wireless modems. Many can hold removable flash memory cards. Some even double as cell phones.

Most PDAs are intended to synchronize with home or office workstations so they usually cannot permanently store data. They may lose any data stored in memory if they lose battery power. Rechargeable batteries typically provide power. Therefore, a battery charger is essential to safeguarding stored data.

Introduction, continued

Other Data Storage Devices

There are many devices that can store and retrieve audio, video, and text data. Each one contains features used to manage data for specific purposes. The following list provides a brief description of each device and the type of data it handles.

Telephones

The three types of phones include cell, cordless, and direct connect to a landline system. They provide communication using landlines, radio transmission, cellular systems, or a combination.

Types of Data Stored: Many phones are programmable and are capable of storing names and phone numbers. Cellular phones can store appointments, e-mail, pages, voice mail, and passwords. Newer cellular phones can have the full features of PDAs.

Phone Answering Machine

An answering machine can be an integral part of the phone or a separate unit that connects the phone to a landline. It records voice messages from callers using either magnetic tape or a digital system.

Types of Data Stored: Phone numbers and names, voice recordings, deleted messages, time/date information, memos, and caller IDs.

Fax Machines

Fax machines transmit and receive documents over the phone system. They have memory capacity to store scanned outgoing documents prior to transmission and incoming pages prior to printing.

Types of Data Stored: Pre-programmed phone numbers, document pages, and a send/receive log.

Digital Cameras

Digital cameras capture images that frequently have associated date and time stamps. These cameras may have built-in memory, which may be expanded using flash ROM cards.

Types of Data Stored: Images in dozens of formats, including any kind of file(s) stored from a computer.

History of Computers

Introduction

One of the first machines to manipulate data dates back to the mid 1600s when Blaise Pascal's Arithmetic Machine automated subtraction and addition computations. Charles Babbage invented the concept of the Analytical Engine that could make decisions for sequential control, branching, and looping based on its own computations. However, Babbage's machine was so massive and complex that he was unable to finish work on it before his death in 1871. These early machines used gears. As electricity was added as a signaling medium, the machines used switches and electro-mechanical relays for computations.

By the mid 1940s, the first electronic computers used vacuum tubes instead of switches. Vacuum tubes could turn on and off much faster than the earlier machines. The vacuum tube computers proved to be very inefficient because they were slow, required large amounts of electricity and space, and generated great amounts of heat.

A noted example of the vacuum tube computer was the ENIAC (Electronic Numerical Integrator Analyzer and Computer) built during World War II by the U.S. Army to calculate artillery and bombing trajectories. This enormous computer contained 18,000 vacuum tubes powered by large amounts of electricity. Even though it could handle 5,000 addition computations a second, one problem could take the staff several days to program. The invention of the transistor would do away with such inefficiencies.

History of Computers, continued


First Transistors

The transistor, a small, solid-state electronic switch, was first invented in 1929 and manufactured in 1947 by Bell Labs. This first semiconductor transistor had no moving parts, was one-fifth the size of the vacuum tube it replaced and one hundred times faster. By the early 1950s, Texas Instruments started producing silicon transistors that paved the way for the small modern computer. During that same period, IBM started selling its Model 650 computer to a few government agencies and commercial businesses.

The Integrated Circuit

The integrated circuit (IC), invented by Texas Instruments in 1959, enhanced computer performance. The first IC contained several transistors and circuitry connected by layers of semiconductor (silicon) material. The connection paths are etched into the silicon “chip” with acid or lasers.

Example of a Modern Printed Circuit Board with Integrated Circuits soldered on



Network Interface Card

As refinements continued, the integrated circuits became miniaturized as many more tiny transistors were placed on a single silicon chip. The microprocessor, developed during the 1970s, can contain several million transistors.

History of Computers, continued

The Growth of Personal Computers

Integrated circuit technology enabled manufacturers to build smaller and cheaper computers. The first personal computer (PC) made its debut in 1975 when Micro Instrumentation and Telemetry Systems produced the Altair 8800. The unit was sold as a kit that contained an Intel 8080 microprocessor and 256 bytes of RAM. These computers were built without a keyboard or monitor. Altair users flipped switches on the front panel to input data and programs. The Altair displayed output on rows of small lights called light-emitting diodes (LEDs).

As the demand for PCs grew, manufacturers devised ways to make the computers user-friendly by adding keyboards, video displays, and data storage devices. The Apple computer, introduced in 1976, was the first PC considered powerful enough to be universally accepted by businesses and average consumers.

Soon, companies like IBM, Radio Shack, and Commodore were offering new products for both business and home use. CPUs became more powerful and offered increased computational abilities. Graphical user interfaces made the new microcomputers user-friendly. Today computer users run multiple programs simultaneously using sophisticated operating systems such as Windows and Unix.

Basic System Components

Introduction

A computer system has a standard set of components that can be divided into four categories.

1. Main system components
2. Data storage and retrieval components
3. Input components
4. Output components

The following sections provide a brief overview of each category. Details for each component will be presented later in this course.

Basic System Components, continued

Main System Components

The main components of a PC include the following devices.

Component	Description
Motherboard	Considered a main component of the computer, the motherboard is a printed circuit board that holds memory chips, expansion cards, and various other components.
Central Processing Unit (CPU)	A device that uses microchip technology to process information and code used by the computer. The CPU is called the brains of the computer.
Bus	A common pathway that data and power signals travel over to various computer components. It is called the computer's nervous system.
Chipset	Main circuit of the motherboard that controls many different components of the system.
Memory	Stores everything a computer system is processing at a given time. Random access memory (RAM) is the computer's short-term memory and the read-only memory (ROM) is the hard-coded memory that is ever-present.
Power Supply	Component that provides power to every piece of hardware within the computer case. It also converts the voltage from a wall outlet to a level a computer can use safely.
Cooling Fans	Fans force air into the case and over components to cool them.
Chassis	The computer's case that houses the system's internal components. Typically, these are constructed in two form factors (shapes): Tower model (approximately 2 ft. long by 10 in. wide, by 2 ft. high) and Desktop model (approximately 2 ft. by 2 ft. wide by 7 in. high).

Basic System Components, continued

Data

Storage/Retrieval Components

Data is stored in and retrieved from the following components:

- **Hard Drive:** Main data repository for non-volatile mass storage. It uses magnetically coated metal, glass or ceramic platters as storage media. Hard drives can be found connected internally in a computer or found within an external enclosure that is attached to a computer through a physical cable or wireless network connection.
- **Floppy Drive:** Portable semi-mass storage that uses 3.5 inch floppy disks.
- **CD-ROM/DVD:** A non-volatile, optical mass storage device.
- **Flash Storage:** A versatile, solid-state storage device that can be used as an additional hard drive or as RAM. They are used in laptops, notebooks, and select PDAs as PC Cards, ExpressCards, or typical USB thumb drives.

Input Components

The following devices are used to input data to the computer:

- **Keyboard:** A primary input device that uses alphanumeric keys.
- **Mouse:** A device that moves a pointer to make selections within a graphical user interface (GUI).
- **Game Controller:** A joystick or other device used to play games. These controllers require a game controller card or sound card, chip, or chipset with a game controller port.

Basic System Components, continued

Output Components

The following devices perform data output:

- **Monitor:** The main display component that interactively shows visual input/output. A monitor requires a video card or video chip/chipset.
- **Video Card, Chip, and Chipset:** Translates visual input/output and sends it to a monitor. These components are found on the MB.
- **Speakers:** Carry audio data processed by the sound card or chipset. They may be attached to the sound card by cables mounted within the chassis and connected directly to the motherboard, PC speaker, or both.
- **Sound Card and Chipset:** Translates audio input/output (I/O) and sends it to the speakers. Both are found on the MB.

This page intentionally left blank.

Lesson 3 – Motherboard and Components

Introduction

The *motherboard* is considered to be the main component of the computer to which all other components are attached. The *Basic Input/Output System (BIOS)* is the instruction set that controls the main functions of the computer. The motherboard holds the ROM chip that contains the BIOS. The *bus* is a circuit that transports data, signals, and power to and from the CPU, memory, and other components on the motherboard. In this lesson you will learn the various components of the motherboard, how a bus works, and the various types of buses found on most systems.

Purpose of this Lesson

You need a basic knowledge of the motherboard and its components to better understand how computer systems function. You will recognize the various types of buses that are found on computer systems to assist in determining the types of devices that can be connected to them.

Objectives

After successfully completing this lesson, you will be able to:

- Define the role of the motherboard
- Identify types of motherboards
- Explain BIOS and the concept of Plug and Play
- Identify main motherboard components
- Know the basic functions of buses
- Identify various bus types
- Recognize various bus connectors

In this Lesson

The following table shows the contents of this lesson.

Topic	See Page
Motherboard Overview	C-18
Motherboard Components	C-23
The Boot Process	C-28
Bus Overview	C-30
Bus Types	C-31

Motherboard Overview

Introduction

The motherboard is the main circuit board of the computer. Every component is connected to it in some way. Motherboards contain slots that hold the processor, expansion cards, and connectors for attaching additional boards. Typically, you will find the following components on the motherboard:

- CPU
- ROM (System BIOS)
- Serial and parallel ports
- Memory
- Chipset
- Clock and Complementary Metal Oxide Semiconductor (CMOS) battery
- Mass storage interface
- Expansion slots
- Connectors for peripherals including monitor, keyboard, and disk drive(s)
- Vista Specific (Screen-duo and ReadyBoost)

Types of Motherboards

There are several different *form factors* (designs) of motherboards. Older form factors include the Baby-AT, which was the first IBM PC board released in 1981, the Full-size AT, and the LPX.

The modern form factors that are found in most computers today include the NLX, BTX and the ATX family of form factors, namely the Micro-ATX, Flex-ATX, and WTX.

Note about PC



The term PC was originally trademarked by the IBM corporation, but today is used widely to mean almost any personal computer. When the term PC is used in this course, it refers to a computer that is based on the Intel X86 processor architecture (discussed later in the book).




Other manufacturers make popular computers that are different from PCs. Apple Computer makes the Mac line of computers, which has experienced dramatic growth in the consumer market. Sun Microsystems computers are popular for government and enterprise business use. Both companies make a full line of products ranging from desktop devices to large servers.

Motherboard Overview, continued

Full-Size Advanced Technology (AT) Board

The full-size AT motherboard replaced the original IBM XT motherboard in 1984. It started out as a large board measuring 12 inches wide by 13.5 inches long, but later was reduced in size as advancements in design progressed. It contains two power supply connectors that plug into one non-form molded power connector and a combination of 16-bit and/or 8-bit ISA slots. These slots, shown in the image below, are used to connect expansion cards to the motherboard. They are discussed in detail in a later lesson.

Example of AT Motherboard




Motherboard Overview, continued

AT Extended (ATX) Board

Intel introduced the ATX in 1996 as a replacement for the Baby-AT. It was considered the first dramatic improvement in motherboard form factors used in desktop PCs. The ATX provided a standard, nonproprietary design that was easy to install and maintain. Many modern motherboards use this same form factor (9 inches wide by 12 inches long.)

ATX was the first to integrate components such as the Flash BIOS and I/O logic. The ATX motherboard is half the width of earlier motherboards and contains combinations of ISA and PCI slots, expansion slots that are covered in depth later in this lesson. The power connector for the ATX is one form-molded power connector that prevents it from being connected incorrectly.

Example of ATX Motherboard



ATX-class and above motherboards may be configured for suspend or power-off functions that are initiated by the operating system. This is especially true with the Microsoft Window 95 or higher versions. During *suspend*, the system goes into a low power state called a sleep mode. Contents of RAM are saved during this state to allow the machine to wake up quickly with all running applications remaining open. The system awakes after the mouse or keyboard is used. During power-off, the system shuts down completely after exiting the operating system.

Motherboard Overview, continued

The ATX Family

The following motherboard form factors were developed by Intel as evolutions of the original ATX:


- Micro ATX was built as a smaller design for use in the first small, low-cost computer systems sold in retail stores for under \$1,000. The reduced size (9.5 inches wide by 9.5 inches long) allowed for a smaller power supply and few I/O bus expansion slots.
- Flex ATX was introduced as a smaller version of the Micro ATX and considered the least expensive motherboard of the ATX family. It will only support the socket-type CPU (size: 9 inches wide by 7.5 inches long).
- WTX was designed as a high-performance ATX. It is a relatively new board used in high-end servers and workstations. It contains a flex slot that is an enlarged PCI slot used to hold powerful multifunction cards (size: 14 inches wide by 16.75 inches long).

Motherboard Overview, continued

Low Profile Extended (LPX) Board

The LPX is a semi-proprietary, non-standard design introduced by Western Digital in 1987. Its low-profile design incorporates slots that are parallel to the motherboard allowing the expansion cards to plug sideways into the riser board. The riser board connects to the motherboard. This design change allowed for slimmer PC cases. LPX was used in PCs sold in retail stores such as Compaq and Packard-Bell. It is easy to identify because devices are parallel to the motherboard. Components for it are difficult to obtain. (size: 9 inches wide by 13 inches long).


Riser Board Example



New Low Profile Extended (NLX) Board

The NLX is a modified, non-proprietary LPX design made by Intel. With the NLX system, the riser plugs into the side of the motherboard. This configuration allows easy access to components for installation and maintenance. The NLX has an integrated network interface card (NIC).

Example of NLX Motherboard with Riser



Motherboard Components

Motherboard BIOS

The *motherboard basic input/output system (BIOS)*, also called the system BIOS, is considered to be the heart of the computer because it controls communications between computer hardware and the operating system. System BIOS is also referred to as ROM BIOS because the code is contained in a non-volatile, read-only memory (ROM) chip. As opposed to typical memory chips, non-volatile memory does not lose its contents when electricity is removed, making the chip suitable for storing data for many years.

The system BIOS contains a software instruction set called firmware. *Firmware* provides the basic input/output instructions to boot the computer and handles several important functions including identifying hardware currently installed in the PC, determining which device will boot the PC, and installing basic drivers for the keyboard, video, and disk drives prior to the operating system loading.

The system BIOS is explained further in depth in Module 3 of this book.

Complimentary Metal Oxide on Semiconductor (CMOS)

CMOS is a chip that stores clock settings, the current system configuration data as discovered by a standard Power-On Self Test (POST) or defined by the setup program, and the Plug and Play settings. Located on the motherboard, CMOS is volatile and requires battery power to maintain the CMOS memory and system time whether the PC is on or off. Battery power comes from one of the following:

- Coin-type watch battery (commonly used)
- Brick/Barrel type battery
- Capacitor, an electrical component that holds a charge

Data stored in the CMOS chip is accessed by the system BIOS and also includes configurable settings such as boot sequence, CPU clock speed, and power management.

Motherboard Components, continued

Chipset

The *chipset* controls the flow of information between various components of the motherboard. The chipset on a modern PC contains two or three separate chips and older PCs had as many as five chips. The largest chip is called the North Bridge; the smaller chip is called the South Bridge. The chipset controls many different components of the system including:

- CPU
- Cache
- Main memory
- Peripheral Component Interconnect (PCI) bus
- Industry Standard Architecture (ISA) bus
- Various system resources

In addition, the chipset defines the various functions the system will support including:

- Defines Front-side Bus (FSB) speed (from 66MHz to over 1000MHz)
- Supports Accelerated Graphics Port (AGP) video cards
- Defines the minimum and maximum processor speed the motherboard can handle

The major chipset manufacturers are Intel, Apollo, VIA, and SIS.

Super I/O Chip


The *Super I/O chip* is the chip on the motherboard that integrates devices that were contained on expansion cards on older PCs. This chip allows for a faster transfer rate of data between the device and the system and has a lower failure rate. The Super I/O chip usually contains the following devices:

- Dual serial port controllers
- Floppy drive controller
- Parallel port controller
- Keyboard and mouse controllers

Motherboard Components, continued

Jumpers

A *jumper* is a small plastic-covered metal clip that is placed over metal pins sticking out of the board. When placed on the pins, the jumper enables electricity to flow to the pins, completing the circuit. A jumper is considered closed when the plastic clip covers the pins.




Use of Jumpers

Use: Jumpers are used to control device settings including processor speed and type, bus speed and CMOS password settings.

Dual Inline Package Switches

Dual Inline Package (DIP) switches are small switches embedded into circuit boards. They are used to configure the system functions including the bus speed, processor speed and processor type. DIP switches are toggled either On/Off or 1/0. Microsoft's standard Plug and Play feature has made the use of DIP switches obsolete.

DIP Switch Example



Motherboard Components, continued

Trusted Platform Module

The TPM is a microcontroller device installed on the motherboard that stores encryption keys, passwords, and digital certificates. It provides secure key generation that can be used to create and/or store both user and platform identity credentials for authentication.

- Offers improved, hardware based security
- Uses RSA and SHA-1 encryption algorithms to create an encryption key for a specific computer
- Encryption keys can be used for full-disk encryption, software licensing, and digital rights management

Power Supply

The computer's *power supply* powers all internal components. Power supplies come in different wattage models ranging from 63.5 to 1000-plus watts. Each unit contains a power transformer that converts voltage from the wall socket to the power level the computer can safely use. The unit transmits a power good signal to the motherboard. This signal must be present continuously for the computer to run. If it is not, the computer shuts down instantly.

The power good signal performs several functions:


- Prevents the computer from starting until the appropriate level of operating voltage is reached
- Interfaces with the computer's reset switch. When the reset switch is pressed, the power good signal is grounded out. When the switch is released, the power good resumes and the system reboots.

Motherboard Components, continued

ATX and AT Power Supplies


When comparing the ATX and the AT power supplies, the main plugs from the power supply to the motherboard are very different. The ATX has a single, form-fitted plastic plug that fits into an on-board socket that has a unique configuration. The plug is form fitted so that it will only fit into the socket in one direction.

ATX On-Board Power Socket (front view)



In contrast, the AT power supply has two separate plugs that fit into two separate on-board sockets. The wires on the plugs are color-coded. To make the correct connection, be sure to place the plugs into the sockets with the black wires of both plugs located directly next to each other. When properly connected, these black wires will be in the center of the two seated plugs.

AT Power Supply Socket (2 cords required)



Warning: Unlike the form-fitted ATX plug, the AT plugs can be connected incorrectly. If the AT plugs are not connected correctly, the motherboard will fail and a fire may occur.

The Boot Process

Overview

All computers are designed to start in a predictable way from the moment you press the Power On button until the moment the operating system loads. The list below outlines the steps.

The Boot Process

Activity	Description
1. Power good signal is sent to the CPU	When you press the Power On button
2. CPU looks at ROM for basic instructions (BIOS)	When CPU receives power good signal
3. System BIOS loads	
4. BIOS initiates Power-On Self Test (POST)	
5. POST checks RAM and then Video. If either of these have a problem, there are various beep codes. <i>From this point forward, errors are reported with text messages displayed on the monitor.</i>	Typically, a procession of long single beeps for RAM; one long and two short for video. Motherboard documentation contains beep codes.
6. When RAM and Video pass the POST test, a single beep occurs. The single beep exists simply to indicate that the diagnostic speaker is working. A malfunctioning speaker will prevent audible beep codes.	You will begin to see text on the screen. The rapid numbers flashing indicate an in-depth RAM check. The screen will indicate the BIOS manufacturer and version number.
7. POST then checks keyboard.	If an error occurs, a text message generally displays the on-screen
8. Legacy and then Plug and Play devices are identified	The data gathered is then stored on the CMOS chip
9. CMOS data is queried against new current configuration data. Drives spin, lights flash, and sounds are heard.	If there is a problem with the CMOS battery, you will get a text message.
10. Finding no major hardware errors, BIOS turns the process over to the boot loader.	

The Boot Process, continued

Overview, continued

Activity	Description
11. The boot loader learns the boot sequence from the BIOS (e.g. A: C: CD-ROM, etc.) and looks for the Master Boot Record (MBR) on that device. <i>For hard disks, the boot loader looks for a partition table. The partition table will have a pointer to the MBR on the primary, active partition</i>	
12. The MBR contains the first file needed to start the operating system (IO.SYS in Windows 9x, boot.ini in NT).	
13. The whole process is turned over to the OS and you see splash screens, etc.	

Bus Overview

Introduction

The various buses comprise the transportation system within every computer. It acts as a highway that sends data, signals, and power among the processor, memory, and other components. In general, there are two bus categories: the internal bus that connects all the internal components to the CPU and main memory and the expansion bus that connects expansion boards to the CPU and main memory.

A computer has several different types of buses. The key buses found in many computers include:

- Processor bus
- Memory bus
- Accelerated Graphics Port (AGP) bus
- Peripheral Component Interconnect (PCI) bus
- PCI Express
- Industry Standard Architecture (ISA) bus
- Universal Serial bus (USB)
- External SATA (E-SATA)

Bus Architecture

Buses are made up of a complex system of thin circuits known as *traces* that are located on any of the several layers of the motherboard. The system chipset orchestrates data transfer from all components via the bus. In addition to circuits, the bus also includes microchips and slots to hold expansion cards or circuit boards.

Buses are hierarchically arranged so that each slower bus is connected to the faster bus above it. The bus size, called width, describes the amount of data (measured in bits) that can be transmitted at one time. The bus's clock speed, measured in MHz, describes the speed of data transfer.

Processor and Memory Buses

The *processor bus* is the data pathway between the CPU and the motherboard chipset. Also called the front side bus, the processor bus is the fastest bus on the motherboard. It is used by the CPU to transfer information between cache or main memory and the chipset.

The *memory bus* is the data pathway between RAM and the CPU. It is always the same width as the processor bus.

Bus Types

ISA Bus

The ISA bus was part of the first IBM PC in 1984. This early version was 8 bits with a speed of 5 MHz. Today, the ISA bus still remains slow at 16 bits and 8 MHz, which is ideal for slow-speed peripherals such as some older modems and sound cards. Until recently, most motherboards contained several ISA slots for backward compatibility. Newer motherboards have replaced ISA slots with a PCI bus.

Extended ISA Bus

The Extended ISA (EISA) bus is a 32-bit, non-proprietary slot connection designed to replace the ISA bus. This bus accepts ISA devices and has two slots. These slots are usually brown in color. The EISA bus is now obsolete in PCs, but they are still used in high-end servers.

Micro Channel Architecture

The Micro Channel Architecture (MCA) connector was an IBM proprietary slot connection designed to replace the ISA/EISA cards. The MCA system is now obsolete, but may still be found in older IBM computer systems. MCA introduced the concept of *busmastering*. This concept allows devices direct access to the CPU via the motherboard I/O controller for faster access. Busmastering is still used by modern devices. There are two formats: 16-bit with two slots and 32-bit with three slots (third slot is separated from the other two slots).

MCA 16-bit bus (2 slots)



MCA 32-bit bus (3 slots)



Bus Types, continued

Video Electronic Standards Assoc. Local Bus

The Video Electronic Standards Association (VESA) local bus (VL-Bus) is a 32-bit, non-proprietary slot connection meant to replace the ISA bus. It has three slots (two together and one separated). The first two slots (ISA) are black and the third parasitic slot is brown. It is used for older video cards and was replaced by the PCI bus.

Example of VL-Bus



Bus Types, continued


PCI and PCI-X Bus

PCI bus is a collection of 32-bit or 64-bit connector slots on the motherboard, generally white in color. Modems, NICs, SCSI host adapters, and non-AGP video cards use the PCI bus. The PCI local bus is also called a mezzanine (meaning intermediary) bus because it sits in the middle between the CPU and RAM. It's part of the North Bridge and can function with other devices and RAM without the use of the CPU.

PCI and PCI-X send data in parallel form in one direction at a time in speeds ranging from 33MHz (PCI) to 533MHz (PCI-X) with a maximum of 34Gbits/sec transfer in the most recent PCI-X. All devices on a PCI bus take turns accessing this bandwidth.

PCI-X is completely backward compatible. All 32-bit PCI cards will function in a PCI-X slot and new 64-bit PCI-X cards will function in a standard PCI slot from the late 1990s. While popular in server environments, PCI-X should not be confused with the newest bus implementation, PCI-Express.

Example of PCI and PCI-X Bus



Bus Types, continued


PCI Express

PCI Express is a PCI advancement that sends data across *lanes* in serial form and is capable of sending and receiving data simultaneously. There are various formats of PCI Express such as x1, x2, x4, x8, x12 and x16. The number, when multiplied by four, represents the number of lanes available to send and receive data. For instance, the x2 has 8 available lanes and the x16 has 64 lanes. In addition to multiple lanes, PCI Express introduces lane switching which allows data to be switched along lanes as needed instead of all devices taking turns waiting for the bus. This makes PCI Express much more efficient than PCI or PCI-X. PCI Express is also known as 3GIO or 3rd generation input/output.

The most common formats are x1 and x2 for general peripheral devices and x16 as an AGP replacement for graphics cards. The x16 format can provide up to 128 Gb/sec throughput.


PCI Express connectors, generally black in color, are physically different from PCI and are not backward compatible with PCI or PCI-X. It is expected that motherboards will contain some combination of PCI and PCI Express for the next few years as the industry makes the transition to all PCI Express.

Examples of PCI Express



Bus Types, continued


PCI Express, continued



AGP Bus

The AGP bus is used exclusively for high-speed graphics processing. This 66 MHz bus has a 32-bit slot connection that is brown in color. It is reserved for a video card. The AGP bus is available in 1x, 2x, 4x and 8x transfer rates. The AGP local bus is placed near the processor bus for direct access to it.

Example of AGP Bus




Bus Types, continued

USB

The USB has a port connection often located at the rear of the computer. Most computers have a USB port that is used to connect various types of peripherals to the computer system. USB brings Plug and Play capabilities to peripherals connected outside the PC. These peripherals are automatically configured when attached to the USB port and a reboot is not necessary to use the component.

Example of USB Ports



Theoretically, you can *daisy chain* (connect various devices to each other in series) 127 devices to each USB port. However, in reality any more than five devices require a USB hub. The current USB specification, version 2.0, supports a data transfer rate of 480 Mb/sec. Version 2.0 is backward compatible with earlier versions 1.1 and 1.0 that used 12 Mb/sec. and 1.5 Mb/sec. rates.

USB is a continually evolving technology, as evidenced by the announcement of USB 3.0. USB 3.0 increases the speed rating of external devices by ten times that of USB 2.0, transferring data at 4.8 Gb/sec. USB 3.0 will be fully compatible with 2.0 and 1.1 devices.

In addition to wired USB solutions, there are recent innovations in Wireless USB (WUSB). WUSB allows for USB 2.0 speeds to devices within three meters of a computer. Wireless USB works similarly to Bluetooth, but features a reduced range to support higher transfer speeds.

Bus Types, continued

Proprietary Bus Connectors

You need to be aware of several proprietary bus connectors. These are explained here.


IEEE 1394 (FireWire)

The IEEE 1394 is an alternative to a USB port. Brand names for this connector are FireWire by Apple and i.Link by Sony. It allows 63 devices to be daisy chained (connected) to each connector. It can also handle multiple chains. The transfer rate is approximately 400 Mb/sec, which is faster than the USB 1.1 but slower than USB 2.0 480 Mb/sec. There is also a 1394b standard that allows faster signaling, up to 3.2 Gb/sec., and requires special cables and interfaces.

1394a (FireWire 400)

1394 and 1394a are the original implementations of this bus type, with speeds up to 400 Mb/sec. 1394a uses two styles of connectors: a 4-pin connector and a 6-pin connector, shown below. The 6-pin has become a standard in many computers and external devices.

IEEE 1394 Connector Example



1394b (FireWire 800)

1394b was designed as a higher performance bus type, allowing up to 800 Mb/sec and greater cable distances. However, 1394b requires a completely different cable and connector, a 9-pin connector that resembles the 4-pin connector shown above. While there are great benefits to 1394b, it has not completely overtaken the original 1394a. Many modern devices feature both 1394a and 1394b connectors.

Firewire S3200

S3200 refers to a newer FireWire standard that uses existing 1394b cables to achieve data transfer rates of 3.2 Gb/sec, four times that of FireWire 800.

Bus Types, continued

E-SATA

The E-SATA is an external port connection for external SATA devices, such as hard drives or DVD/CD devices. E-SATA runs at speeds of 300 Mb/sec.

Example of E-SATA Cable and Ports



Lesson 4 – CPU and Memory

Introduction

The computer's processor, called the central processing unit (CPU), works in concert with memory to process software and user commands. This lesson explains the significance and functions of both CPU and memory.

Purpose of this Lesson

You should understand how a computer processes commands and data. You will be presented with information about CPUs and memory to enhance your general knowledge of computer systems.

Objectives

After successfully completing this lesson, you will be able to:

- Explain the basic functions of the CPU
- Identify the CPU in a computer
- Recognize various types of memory

In this Lesson

The following table shows the contents of this lesson.

Topic	See Page
CPU Functions	C-40
Memory	C-41

CPU Functions

CPU Defined

The computer processor, or CPU, is the main component that processes all software instructions and makes all calculations.

The CPU's main components are the following:

- Arithmetic logic unit that handles all arithmetic and logical operations
- Control unit that takes instructions from memory, translates them, and then carries out the instructions.

The CPU can be either a single-socket chip or a chip that is mounted on a slot circuit board.

Types of CPUs

There are various types of CPUs from several manufacturers. Each type connects to the motherboard in different ways. Some CPUs connect via a socket in the motherboard and others use a slot connector. Each style of connector is produced in a variety of specific and unique designs, few of which are compatible with each other. A CPU and a motherboard must be completely compatible for a computer to operate.

Socket-type CPUs were the original design and is still the most popular. The actual CPU chip attaches to the motherboard through a pin-grid array (PGA), a square receiver containing hundreds of evenly spaced holes. Modern Intel CPUs are designed as Plastic PGA (PPGA) or Flip-Chip PGA (FCPGA). The only physical difference between these is the orientation of the actual processor chip in relation to the motherboard.

During the late 1990's many CPUs were designed using a slot-type connector. Similar to a PCI card, the CPU would seat into a rectangular slot on the motherboard. However, this design was inefficient for faster clock speeds and the design was phased out. The Pentium II and many Pentium III processors were designed as slot CPUs.

Memory

Introduction

Memory is the temporary data storage area of a computer system. It is the system's workspace that houses the programs and data being processed by the CPU. The two main types of memory are:

- Read-only memory (ROM) is non-volatile and cannot be written to. The ROM chip contains the motherboard BIOS that is used to boot the system.
- Random access memory (RAM) is the main memory that can be read by the CPU and written to. RAM is temporary because it relies on electrical power. RAM is also referred to as physical memory that is represented by the actual computer chips that hold data.

Memory is often confused with disk storage (on hard drives, disks, or tapes) because both are measured in megabytes and gigabytes. To remember the distinctions between memory and disk storage, consider this analogy. RAM represents your current work files on your desk that are easily accessible when you want to make changes to them.

Conversely, disk storage represents an area where you permanently store completed work files, such as in a file drawer. You access those files less frequently and you must do a search to retrieve them.

Memory, continued

ROM

There are four types of ROM chips as detailed below. All types use non-volatile data storage meaning that the data remains indefinitely on the chip until the chip is reprogrammed using special hardware or software.

- **ROM:** During manufacturing, binary data is stored in the die of the silicon and cannot be changed without making a new chip.
- **Programmable ROM (PROM):** This chip comes blank and needs to be programmed using a special machine called a device programmer. Once the PROM is written to, it cannot be changed.
- **Erasable PROM (EPROM)** – A type of ROM that can be erased by exposing it to high-intensity ultraviolet light. The die used is sensitive to ultraviolet light. When exposed, all the binary 0s are changed back to 1s.
- **Electrically Erasable PROM (EEPROM) or FLASH ROM:** Chips that can be electrically erased and reprogrammed on the circuit board using a special software program. No other special equipment is required to reprogram the chip. All modern MBs use this type of chip for the system board BIOS.

RAM

RAM is short-term memory used to store data being processed by the CPU. RAM is volatile because any data that is stored in RAM is cleared when the power is off or the system is reset. RAM chips are mounted on sticks that fit into connection slots on the MB. Slots are arranged in numbered banks starting at 0. Bank 0 is usually located near the CPU.

Memory, continued

Dynamic RAM (DRAM)

DRAM is the most common and least expensive memory chip. It is small and has high data density (up to 256K). Other characteristics include:

- Several DRAM chips are mounted on a single stick that fits into connection slots on the MB.
- Connection slots are arranged in banks starting at 0. Bank 0 is usually located near the CPU.

DRAM requires constant electrical refreshing to keep it dynamic. This is done by using capacitors and transistors in pairs. Capacitors hold charges (both positive and negative) that indicate whether the transistor is On or Off. The charge holds power in the transistor and keeps the RAM contents alive.


Memory, continued

Single Inline Memory Modules (SIMM)

SIMM sticks were used in earlier generation PCs. Manufactured in sizes ranging from 512 KB to 32 MB, they were manufactured with flat connection pins in two types:

1. 30-pin SIMM, considered obsolete in PCs (used with 386 processors and below), is installed in pairs
2. 72-pin SIMM is still available and is installed in singles on motherboards with CPUs below Pentiums and in pairs on Pentiums and above. These were widely used a few years ago until they were replaced by DIMMs.

Example of SIMM Chip



A SIMM has a unique insertion method and locking mechanism. SIMMS are inserted into the connector slot at a 45-degree angle and then rotated into the slot. The locking mechanisms are small metal or plastic pins located on the side of the slots. You know the module is a SIMM if you remove it at a 45-degree angle.

Memory, continued

Dual Inline Memory Modules (DIMM)

The most widely used memory modules are the 168-pin DIMMs, which have replaced SIMMs. They are 64 bits wide and range in size from 8 MB to 1 GB. DIMMS are inserted straight into their sockets and locked in place.

Synchronous DRAM (SDRAM)

SDRAM is the modern memory standard and is faster than regular DRAM. It is most prevalent in Pentium II/Athlon systems and above. SDRAM runs in synchronization with the actions of the processor bus. It performs operation at the same time as the system clock and at the same speed. This increases the speed of the data input/output. All SDRAM are DIMMs. All DIMMs are not necessarily SDRAM.

SDRAM is manufactured at single, double data rate, or quad data rate (SDR, DDR, DDR2 and DDR3). These rates describe the speed of data transfer per clock cycle. There are several SDRAM speeds. Which type and speed your system requires depends on the motherboard.

The following lists various SDRAM speeds and their associated processor bus speeds:

SDRAM Speed	Associated Processor Bus Speed
PC 66	66 MHz
PC 100	100 MHz
PC 133	133 MHz
PC 2700	333 MHz
PC 3200	400 MHz
PC 4200	533 MHz
PC 5300	667 MHz
PC 6400	800 MHz
PC 10666	1333 MHz
PC2 8500	1066 MHz
PC3 12800	1600 MHz
PC3 14900	1866 MHz

Memory, continued

Synchronous DRAM (SDRAM), continued

When upgrading SDRAM, it is important to select speeds that are greater than or equal to the frontside (processor) bus speed. Pushing slower RAM at higher speeds (overclocking) will cause the RAM to overheat and malfunction, which can also cause permanent damage to the motherboard. It is acceptable to install faster RAM into a slower motherboard. SDRAM is not compatible with all motherboards, so it is important to refer to the motherboard documentation.

RIMM (RAMBUS Inline Memory Modules)

RIMMs are found in high-end computers because they offer the highest performance of available memory, with transfer speeds capable of over 6 GB/s. The memory modules use RAMBUS Dynamic RAM (RDRAM) chips, a proprietary chip format. These are generally geared towards the server market, but Intel actively pushed the technology towards the consumer market. Due to their high cost, they are uncommon on most computers.

Example of RIMM Chip



RIMMS are manufactured with 184 connection pins and in sizes that range from 64 MB to 1 GB. Most motherboards require that RIMM chips be installed in pairs. If only a single memory module is needed, then a special continuity unit (CU) is required to provide termination. CUs are additional RIMM modules without the RDRAM chips that are plugged into the remaining RIMM slots not occupied by memory modules.

Module D

Data Storage Components

Overview Understanding the vast array of data storage components is vital knowledge for processing an electronic crime scene investigation. This module introduces disk drives and various types of removable storage media.

Purpose of this Module The purpose of this module is to introduce students to the various types of media available for the storage of digital data.

Objectives After successfully completing this module, you will be able to:

- Explain how data is stored on a hard drive
- Identify components of the hard drive
- Discuss the workings of a floppy drive
- Recognize various removable media

In this Module Here are the lessons in this module:

Lesson	See Page
Lesson 1 – Hard Disk Drives	D-3
Lesson 2 – Floppy Drives and Removable Media	D-35

This page intentionally left blank.

Lesson 1 – Hard Disk Drives

Introduction

Hard drives are the main storage of the computer. Drives use highly sophisticated technology to write data on platters. This lesson examines the main components of hard drives and their functions. It also introduces you to data storage methods.

Purpose of this Lesson

In this lesson, you will learn how disk drives store information. This is important to knowing how to safeguard data during a crime investigation.

Objectives

After successfully completing this lesson, you will be able to:

- Identify the main components of a hard drive
- Explain the process by which data is stored on and retrieved from a hard drive
- Describe the basic formatting procedures for hard drives
- Explain hard drive geometry

In this Lesson

The following table shows the contents of this lesson:

Topic	See Page
Hard Drive Overview	D-4
Hard Drive Components	D-5
Hard Drive Controllers	D-11
Hard Drive Geometry	D-17
RAID	D-19
Drive Preparation	D-22
Hard Drive Preparation	D-24

Hard Drive Overview

Saving a File to a Hard Drive

To understand how the hard drive works, you first need to know how a file is saved to it. As a file is being written or created, all of its contents are temporarily stored within RAM. When you save the file, the software program you are using makes a request to the operating system to take the file from RAM and store it permanently to the hard drive. Your request to save a file initiates a complex process that records your file and tracks its storage location on the hard drive.

Process for Storing Data on a Hard Drive

For illustration purposes, the following table describes the general process for saving a file to a hard drive in modern operating systems.

Step	Activity
1	The user makes a request to save a file. That file is then temporarily stored in RAM (if not already residing there).
2	The file system analyzes the disk to find the required available space.
3	The operating system receives the request to transfer the file from RAM to permanent storage on disk.
4	The file system (FAT, VFAT, FAT32, NTFS, ext2 or ext3) updates the file directory with the newly saved file name and its exact location or directory path.
5	The operating system tells the drive controller to save the file. The file is then transferred from RAM to the hard drive.
6	Once the file is saved, the file system marks that area of the disk as not available. That area cannot be overwritten by subsequent file saves.

Hard Drive Components

Components

Hard drives contain the following components:

- Disk platters
- Read/write heads
- Head actuator
- Spindle motor
- Jumpers and/or switches
- Disk controller
- Cables and connections


Disk Platters

A hard drive stores archival copies of all programs and data on non-volatile disk platters made of metal or glass with a magnetic medium coating. Most hard drives can hold several 3.5-inch platters with a current capacity of 250 GB or higher. Laptops and notebook computers use 2.5-inch platters with a current capacity of 100 GB each. Data is stored on both sides of each disk platter.

When in operation, the disks spin and the read/write heads move over the disks and store information in tracks and sectors. Data is stored in concentric rings or tracks on the disks. The tracks are divided up into segments called sectors. Each sector can store approximately 512 bytes of data.

Hard Drive Components, continued

Read-Write Heads *Read-write heads* are the mechanisms that store and read data on platters. There is one read-write head combination for each side of the disk platter and all heads are mounted on a single rack. Heads move in unison across the platters. They float above the surface of the disk platter on a cushion of air generated by the action of the spinning disk platter. The heads float three to five millionths of an inch above the platters. As they move, the heads read changes in the disk's magnetic coating, called magnetic flux. This is the process for interpreting stored data.



Data is stored on both sides of the platters. Side numbering starts at zero for the top side of the top platter. For example, if the drive has four platters (eight sides), the numbering for the sides would be zero through seven.

Hard Drive Components, continued

Jumpers

A jumper is a set of pins used to control device settings. The pins act like an on/off switch and are configured by placement of a small plastic block, called a shunt. When there is no shunt over the jumper pins, the circuit is open, or off. When there is a shunt placed over the pins, a small wire inside the shunt connects the pins and the circuit is closed, or on. This is also called shorted. A shunt can also be attached to a single pin in a parked position. This does not close the circuit; the parked position is simply used to store shunts that are not currently needed.

You use jumpers when installing PATA (Parallel ATA) IDE devices. Typical motherboards house two PATA IDE channels that are identified by a Primary (IDE 1) Connector and a Secondary (IDE 2) Connector. Each IDE channel can support a maximum of two IDE devices. The relationship of the two devices on a single channel is master and slave, which simply designates the sequential order of the two devices. The four possible IDE device relationships are:

- Primary Master: Device 1 on the primary IDE 1 channel
- Primary Slave: Device 2 on the primary IDE 1 channel
- Secondary Master: Device 1 on the secondary IDE 2 channel
- Secondary Slave: Device 2 on the secondary IDE 2 channel

The designation of master and slave between two devices on the same IDE channel is determined by jumper settings on the IDE devices. The drive normally contains a diagram indicating which pins to short and which to leave open. The purpose of setting drives on one connector or another, and choosing between master and slave, is to place the drives within a specific order. On modern computers, drives are read by the operating system in order from primary master down to secondary slave. Therefore, a drive placed earlier in the chain will have a lower drive letter.


Some motherboards support cable select (CS) technology. By setting the drive's jumpers to CS and using a special CS cable, you can control the master and slave designations of the drives' placement on the cable. Standard master and slave settings as described above are used more frequently than CS.

Hard Drive Components, continued

Jumpers, continued **Setting the Master Drive**

To indicate a primary hard drive in a series of two, close the Master jumper for the following configuration:


Master



Setting the Slave Drive

To indicate a secondary hard drive in a series of two, close the Slave jumper for the following configuration:

Slave




Note that on newer motherboards, with support for SATA (Serial ATA) devices, the motherboard may have only one PATA connector or none at all. Additionally, SATA devices do not require a jumper for Master and Slave settings, as each device is plugged onto its own motherboard connector.

Hard Drive Components, continued

PATA Data Cables Older PATA IDE hard drives and CD-ROMs connect to the motherboard via a ribbon cable to facilitate data transfer. These cables are flat and wide with wires running in parallel the length of the cable. IDE ribbon cables generally have three 40-pin connectors: one to attach to the motherboard and one each for Master and Slave device. A red or black stripe along one edge of the cable indicates the location of Pin 1 and is used to determine cable orientation.

40-pin Ribbon Cable



Connections on devices such as hard drives and the motherboard are usually notched or fitted so that the ribbon cable can only attach one way. If not connected correctly, the device will not work and damage may occur. Always attach the ribbon cable on IDE devices with the red (or black) stripe closest to the power connection.


Hard Drive Components, continued

SATA Data Cables SATA (Serial ATA) is a communication bus technology that is replacing older PATA (IDE) technology. SATA, unlike PATA, does not use Master and Slave jumpers on the hardware to determine device priorities. Instead, each SATA device has its own dedicated connection to the host device. Therefore, no bandwidth is shared with any other devices over the SATA data cables.

SATA offers other benefits over PATA, such as more compact cables for better airflow, and hot swapping capabilities.

SATA data cables use 7-pin connectors that only utilize four wires for transferring data. The other three wires are used as ground.

SATA Data Cable



A new power connector is also specified by the SATA standard, although many SATA devices include both SATA and Molex power connectors. This cable is designed with 15-pins and supports three different voltages: 3.3 V, 5 V, and 12 V. Nine of the 15 pins are used for power, five are used for ground, and the last pin is used for staggered spinup. Staggered spinup allows the drives to initialize and power up sequentially to increase reliability and prevent power surges.

SATA Power Cable



Hard Drive Controllers

Drive Controllers Computer storage devices including hard drives and floppy drives require controllers to govern how they operate. A *controller* is a device that handles the transfer of data between the component and the computer. The disk drive controllers that will be discussed in this course include the following:

- Integrated Drive Electronics (IDE)
- Enhanced IDE (EIDE)
- Serial ATA (SATA)
- Small Computer System Interface (SCSI)
- Serial Attached SCSI (SAS)

Integrated Drive Electronics (IDE) Controller


Integrated Drive Electronics (IDE) is a term used to describe any disk drive with a built-in controller. The technical name for IDE is Advanced Technology Attachment (ATA) IDE. The IDE is the primary interface electronics (controller) that connects a hard disk drive to the computer.

With today's technology, IDE is integrated into the drive. The drive attaches to a connector on the MB. IDE makes the drive more reliable than drives that have separate ISA slotted controllers. This reliability enhances the integrity of data.

IDE controllers are customized to fit the IDE drive. In contrast, separate drive controller cards are generic and may not provide full speed or functionality. IDE controllers contain an independent BIOS that limits the total hard drive size to 528 MB. It uses a standard 40-pin ribbon cable to connect the drive to a motherboard.

IDE Controller Connection

Rear View



Note: If you seize a drive attached to a legacy controller card, be sure to take the controller card as well as the drive to ensure that you can access the disk.

Hard Drive Controllers, continued

Enhanced IDE Controller

Enhanced IDE (EIDE) controllers offer enhanced controller BIOS that increases maximum drive capacity to more than 528 MB. This controller's technical names are Advanced Technology Attachment series ATA-2 through ATA-6. It is marketed as Fast ATA and Fast ATA2. All ATA series drives are backward compatible with older models.

EIDE provides two data channels per connector; therefore, two drives can be connected to each EIDE port. In addition, EIDE allows attachment of Advance Technology Attachment Packet Interface (ATAPI) devices, including CD-ROM, DVD, and Zip drives.

ATA-1, 2, and 3 Controllers

- Use the CPU for data transfer (PIO mode) which takes CPU transferring time away from other tasks
- Data transfer rates: ATA-1 is 8 MB/sec; ATA 2 and ATA-3 are 16 MB/sec

ATA-4, 5 and 6 Controllers

- Uses Direct Memory Addressing (DMA) modes for data transfer. During DMA transfer, the CPU is not used.
- ATA-4 support a transfer rate of up to 33 MB/sec
- ATA-5 supports transfer rates of up to 66 MB/sec
- ATA-6 supports transfers rates of up to 100 MB/sec
- To achieve rates of speed of 66 MB/sec or more, a special 40-pin, 80-wire ribbon cable must connect the hard drive to the motherboard.

PATA Naming Convention

While ATA, IDE, and EIDE technically refer to three separate concepts, they are generally grouped together and referred to as the same technology. Many times, the terms are used interchangeably to refer to hard drives, with the exceptions of SCSI and SATA. Since the introduction of SATA (Serial ATA), all examples of prior hard drives have been retroactively renamed to PATA (Parallel ATA).

Hard Drive Controllers, continued

Serial ATA (SATA)

The two-inch wide parallel PATA ribbon cable has reached a maximum transfer rate of 133 Mbps (megabytes per second), marking its technical limitations. As a need grew for higher performance hard disks, the Serial ATA (SATA) standard was introduced and implemented. SATA currently has speeds of 150 Mbps (SATA 1.5G) and 300 Mbps (SATA 3G) and has the potential to go up to 600 Mbps (SATA 6G). Serial ATA's power requirements, 250 millivolt versus PATA's 5 volt, have made it widely adopted in new low-power motherboards.

Serial ATA drives are connected via a .25-inch cable that connects the drive to a Serial ATA card plugged into a PCI slot or a slot integrated onto the motherboard. The SATA cables have seven pins and seven wires. The thinner cables permit better airflow and are smaller and easier to route. Parallel ATA cables are two inches wide and have a maximum length limitation of 18 inches. The SATA cable has a maximum length limitation of one meter.

Each SATA connection supports a single drive. As previously mentioned, this means that you no longer set jumpers for the master or slave configuration. The SATA standard also has hot swapping designed into its configuration which allows you to swap drives while the system is running.

SATA hard drives also feature a unique SATA power connector. Unlike the standard Molex 4-pin connector typically found in PCs, the SATA power connector features 15 pins to provide 3.3 volt, 5 volt, and 12 volt power, depending on a device's needs. Many current SATA hard drives support both SATA Molex power connectors; however, the SATA power connector is required to take advantage of hotswapping.

SATA channels and drives are not inherently backward compatible with IDE technology, although there are SATA-to-IDE adapters available.

Hard Drive Controllers, continued

External SATA (eSATA)

To handle the growing market of external storage devices, eSATA was designed to provide high-performance data transfers to portable devices. The external SATA connection allows for SATA drives to be connected externally with transfer speeds up to 300 Mbps, which is six times faster than USB 2.0. The eSATA shielded cables are found in lengths up to two meters, but do not provide power to the end device. While eSATA is currently geared towards external SATA hard drives, it is able to accommodate a large variety of external devices, including CD-ROMs.

Small Computer System Interface (SCSI)

SCSI is a system level interface that enables the connection of various peripheral devices to the computer. Most modern home PCs do not come with SCSI hardware preinstalled. However, any computer with a PCI slot can become SCSI compatible.

SCSI is not a controller like IDE. Instead, it is a separate data bus that is connected to the system bus via a host adapter. There are some high-end computers that have a SCSI adapter card or an adapter built into the MB.

Many current SCSI busses can support between 8 and 16 devices, although support for one device is lost to the host adapter. Other SCSI implementations can support many more devices, such as SSA which supports 96 and SAS (Serial Attached SCSI) which supports 16,256. Devices are strung together in a chain and each device is assigned a SCSI ID. A typical SCSI host adapter can support a maximum of 16 devices, although the actual SCSI adapter counts as one of those 16 devices. When one device wants to communicate with the system bus, the data passes through the host adapter to the system bus.


Because the SCSI bus operates like a chain, it must have termination at the end the chain. Most computers can support up to four host adapters.

A SCSI bus can be either 8-bit or 16-bit. The 16-bit bus is typically named wide SCSI.

Hard Drive Controllers, continued

Example: SCSI Card and Cable

Example of SCSI Card and Cable



Serial Attached SCSI (SAS)

Serial Attached SCSI (SAS) is another serial bus technology with similarities to both SATA and SCSI. SAS uses the same communication protocols as SCSI to transmit data and uses the same type of connection cables as SATA. SAS is primarily used in corporate and enterprise environments.

An SAS controller is backwards compatible with SATA devices. Therefore, a SATA hard drive can be connected to an SAS controller and function properly. This allows less expensive SATA drives to be utilized on SAS systems while providing the upgrade path to the higher end SAS hard drives.

SATA controllers do not support SAS devices.

Solid State Drives (SSD)

A Solid State Drive (SSD) is a storage device for notebooks and desktops that uses solid state memory to store data. There are no moving parts in a solid state drive. This eliminates a lot of the seek time and latency along with other electro-mechanical delays that were attributed to a conventional hard drive. Sizes currently range from 32 GB to 256 GB; however, SSD drives are still considered too expensive to be in widespread use.

Hard Drive Controllers, continued

Hybrid Hard Drive (HHD)

Recently there has been much development into hybrid drives that combine solid state and magnetic disk drives. The solid state component acts as a cache for the hard drive and can be up to 1 GB in size. As data is saved from the computer it is temporarily written to the SSD cache. Only when the cache is nearly full, or when new data must be read, does the magnetic portion of the drive spin up. This implementation greatly improves hard drive performance and reliability, while reducing heat and noise generated by normal hard drives.

The flash memory, or solid state portion of the drive, could also be used with Windows Vista's ReadyBoost to increase performance on compatible Vista systems. As of this time, HHDs are only compatible with Windows Vista-based systems.


Hard Drive Geometry

Hard Drive Geometry

Hard drives are divided into tracks, sectors, heads, and cylinders. Each one is detailed here.

Most motherboard BIOS chips prior to 1997 do not automatically detect geometry information. This information is usually annotated on a sticker fixed to the drive. If not, check user reference manuals for the manufacturer's Web site for a listing of this information by hard drive type. It is a good idea to record these values for future reference.

In computer forensic investigations, it is important to know where data is stored on a hard disk because you may be asked during trial about the exact location of a hidden data file. For example, you may be asked to identify the track, sector, and cylinder on which a hidden file is located on a seized computer system.



Tracks

Tracks are the concentric circular paths that are placed on both sides of the platter. Tracks are a specified area that the read/write head hovers over. They are arranged starting at the center of the platter working to the outer edge.


Tracks are measured in density called tracks per inch (TPI). The first PC hard disks in 1982 had 200-300 TPI. Today, drives have over 100,000 TPI. Tracks are uniform on every platter in the drive and a *cylinder* is formed by alignment of the same tracks on all platters.

Hard Drive Geometry, continued

Sectors Sectors are shaped segments of tracks. Each sector holds 512 bytes of data. With today's technology, each track can have between 17 and 100 or more sectors.

Heads Heads on the platter represent the side of the platter. These are sometimes confused with the read/write heads, which are the devices that relay data to and from the platters.

Cylinders All platters in a stack are aligned with each other and they move in unison. A cylinder is formed by identically positioned tracks on every platter.



Clusters A *cluster* is a group of one or more sectors that form the smallest addressable area of storage on a disk. Although a cluster is the smallest unit of disk space used by an operating system, cluster sizes vary and depend upon the OS and the partition size of the disk.

RAID

RAID

A Redundant Array of Independent Disks (RAID) is a method of combining multiple hard drives and storing data in different locations on those drives at the same time. Instead of writing data to individual drives, such as drive C and drive D, a RAID allows the OS and the user to view and use drives C and D as one logical drive. Using multiple drives that appear as one logical drive instead of using one large hard drive allows for greater performance, capacity, and reliability.

There are two methods of writing to the RAID: striping and mirroring. Various types of RAID are produced using different combinations of striping, mirroring, error correction, and parity. RAID is created using software or hardware methods. The most common types of RAID include:

- RAID 0
- RAID 1
- RAID 3
- RAID 4
- RAID 5
- RAID 6
- RAID 0+1
- RAID 1+0

Striping

Striping involves partitioning each drive's storage space into units ranging from 512 bytes to several megabytes. Data is written in bytes or groups of bytes across multiple drives as specified by the interleave ratio. Since more than one drive is reading and writing data at the same time, performance is greatly enhanced.

Mirroring

Disk *mirroring* is a technique that stores the same data on a pair of disks. Mirroring ensures that you always have an exact duplicate of the drive providing fault tolerance.

RAID, continued

Parity

Parity is a method of checking data when it is copied from one storage place to another to ensure that the data has not been lost, overwritten or corrupted. When a group of bits are copied, an additional parity bit (binary digit) is added. This bit ensures that the data has been copied successfully.

In RAID, parity is used for data recovery. It is implemented through *exclusive OR (XOR)*, a logical process that returns a value of “1” if two data bits are not the same. If any two bits are the same (both 1s or both 0s), the result is 0. If they are 1 and 0, the result is a 1.

This XOR process is shown in the following table:

Data Bit A	Data Bit B	Output
0	0	0
0	1	1
1	0	1
1	1	0

You then compare each data bit in a data set to get the XOR result or “parity data.” Here’s an example of how parity data is calculated:

Data A = 10100101
Data B = 11110000
Parity Data = 01010101

Notice that each data bit is put through the XOR process, which is shown in the table. The first digits of Data A and B are both 1s. They are alike. Therefore, the result is a 0 for the first digit of the XOR result. For the second digits, Data A and B do not have similar binary numbers. The 0 and the 1 yield an XOR result of 1. This process continues for each bit in a data set.

In this example, Data A, B, and the parity data are all on separate drives in the RAID. If any one drive becomes missing or corrupt due to failure, you can reconstruct the missing data using the remaining two drives through the same XOR process.

RAID, continued

ECC

Error Correction Code or Error Checking and Correcting (ECC) looks for errors in data that is being read or transmitted. Unlike parity that resends faulty information, ECC attempts to correct the information.

ECC uses a mathematical code to describe each 64-bit word and sends that code with the data. When the data is about to be stored, the code is regenerated. If there is a match, the data is saved. If the codes are different, the original code is used to correct the bits.

Types of RAIDs

The various types of RAIDs work in different ways. Here are the five most common types.

RAID 0

RAID 0 implements a striped disk array without any mirroring, ECC, or parity. The data is simply placed across several drives allowing each drive to read and write data at the same time. This configuration provides the best efficiency and performance, but provides no fault tolerance.

RAID 1

RAID 1, also known as disk mirroring, writes the same data to a pair of hard drives. This implementation provides for the best fault tolerance, but the most overhead.

RAID 3

RAID 3 implements striping small amounts of data across several hard drives with one drive assigned to store parity information. Some RAID controller card manufacturers do not support this type of architecture because a RAID 5 with small stripes yields similar results.

RAID 5

RAID 5 implements striping large amounts of data across several hard drives while rotating parity throughout all the attached drives. This array requires a minimum of three drives.

RAID 0+1

RAID 0+1 implements a mirrored array of RAID 0 drives. This type offers excellent performance and fault tolerance. A minimum of four drives is required.

Drive Preparation

Data Allocation and Storage

Along with all of the differences between drive controllers and disk connectors, there are also differences in the ways that data is stored onto hard drives. As information is stored onto a hard drive platter, it must be placed into an addressable location for later recollection.


Logical Block Addressing

Logical block addressing (LBA) is a process that assigns linear numbers to all sectors in a drive. This process, now a standard, extends the drive's storage capacity from the normal 528 MB. Older motherboard BIOS types or operating systems may require LBA to be translated into CHS geometry. If required, this translation is performed automatically. The 28-bit LBA supports a partition as large as 137 gigabytes. The newer 48-bit LBA will theoretically support a single drive with storage of 144 petabytes.

Zoned Bit Recording

Zoned bit recording lessens wasted storage space on the outer sections of the hard drive by reassigning the sectors on a disk according to the physical size of the track. Data in small tracks near the center of the disk is written closer together than data that is in the larger outer tracks. Therefore, outermost tracks can contain more sectors per track than the smaller inner tracks. The drive controller can recognize the variable number of sectors per track in different zones. Zoned bit recording is set during the low-level format.

Hard Drive with Zoned Bit Recording



Drive Preparation, continued

Perpendicular Recording

Perpendicular recording is a technology that increases the amount of storage space on a hard drive. Using perpendicular recording, manufacturers have produced micro-drives with over 100 gigabytes of storage on a 1.8" drive, a 320 gigabyte 2.5" laptop drive, and a 3.5" drives with over 1 terabyte of storage. This recording method can increase the storage density of a hard drive by almost ten times compared to longitudinal recording.

Longitudinal recording, which has been used since the inception of hard drives, stores bits side-by-side on a magnetic surface. Traditionally, when manufacturers want to increase the amount of storage space, they decrease the size of the magnetic particles on the hard disk surface. This frees more storage space on the disk. However, if the magnetic particles become too small, the data is susceptible to corruption. The risk of data corruption limits the amount of storage capacity with longitudinal recording.

Perpendicular recording vertically records the bits of a sector to a hard drive. Stacking the bits vertically increases the amount of storage density. With perpendicular recording, more data can be stored with less risk of corruption.

Hard Drive Preparation

Hard Drive Formatting

A hard drive must be prepared before using it to store data. This preparation process involves two *formatting* steps (low-level formatting and high-level formatting) with a *partitioning* step between them. Here are the basic steps to prepare a hard drive:

Step 1: Low-level Formatting

The purpose of low-leveling formatting, also called physical formatting, is to set the drive geometry by first setting the tracks and then dividing them into sectors.

Step 2: Partitioning

To use a hard drive, it must be partitioned. To understand partitioning, consider this analogy:

A farmer owns acres of fields. He can both plow and plant the entire parcel at once or he can separate the land into smaller plots and seed them independently at different times. The same is true for a hard drive. You can partition all of the space as one large parcel to hold all of your operating system and data. Or you can partition the drive into smaller plots and use each one independently. In fact, after each partition is formatted, it will appear to the operating system as a separate drive.

When you purchase a computer, the original equipment manufacturer (OEM) has partitioned the drive for you as a single partition. This partition shows up on your system as the C: drive.

However, partitioning your drive into multiple drives offers many benefits. For instance, you could place your operating system on one partition and all of your data on another. If you need to reinstall the operating system, you could do so without affecting your data partition. In addition, you can format different partitions with different file systems in order to install a different operating system on the same physical drive.

There are many ways to partition a drive and the operating systems provide the tools. MS-DOS uses the FDISK utility. Windows XP uses Disk Management and some distributions of Linux use Disk Druid. While the interfaces look different, the result is the same.

Hard Drive Preparation, continued

Step 2: Partitioning, continued

There are only two types of partitions: primary and extended. You need at least one primary partition to boot a computer from a hard drive. Extended partitions are typically used for storing data and applications and cannot be used to boot a computer.

Hard drives can support up to four primary partitions or three primary partitions and one extended partition. However, the FDISK utility only recognizes one primary partition per physical drive, so DOS and Windows 9x/ME can only have one primary partition per drive. Linux, Unix, Windows 2000, XP and Vista can all support four primary partitions.

Once a primary partition is created, it is ready for high-level formatting. Extended partitions, however, require one additional step. Before an extended partition can be formatted, it must be further organized into logical drives. This can be accomplished with the same partitioning utilities mentioned earlier. Once you create the partitions and logical drives, you can proceed to the next step – High-level Formatting.

Step 3: High-level Formatting

High-level formatting, also called *logical formatting*, is performed by the operating system. High-level formatting writes the necessary operating system-dependant file structures to the drive so that the drive can manage and store data. This includes establishing the boot record, file allocation table, and the root directory. The high-level formatting process also scans the disk for areas that are unable to store data. These areas are marked as bad sectors and are isolated.

MS-DOS FAT

MS-DOS, Windows 9x, and Windows ME all make use of the MS-DOS-based FDISK utility. FDISK is responsible for the creation and deletion of partitions and logical drives. Once these partitions are created, they cannot be altered easily. File Allocation Table (FAT) partitions cannot be resized without the use of third-party software. Rather, the original partition must be deleted and then the new partitions can be created.

Warning: Any data residing on a partition will be destroyed as soon as the partition is removed. Any data that is to be saved from the partition *must* be backed up before removing the partition.

Hard Drive Preparation, continued

FAT File System Characteristics

Each partition or logical drive contains a FAT. The FAT manages a list of where files begin and end on the particular drive. This list is constantly being updated with new information as files are saved, deleted, renamed or moved. It is critical that this list stay up-to-date and intact to allow the operating system to access files. It is considered so critical that the latest backup copy of the FAT is saved in another area on the hard drive. This backup copy is used to restore the original if it is damaged or altered.

FAT16 used a 16-bit length number to identify the various clusters on a hard drive. This limited the addressable capacity of a hard drive to 65,526 clusters, or 2 gigabytes of data. Those using FAT16 on drives that were larger than 2 Gb were required to partition the drive to fully use the drives capacity.

FAT32 uses a 32-bit length number to identify all of the clusters on a hard drive. This increase allows the file system to address approximately 268,400,000 clusters, or approximately 2 terabytes of information.

Hard Drive Preparation, continued

FAT Drive Identification

FAT 16/32 systems can support up to two partition types per hard drive. Those types are primary and extended. The extended partition contains logical drives.

Each primary partition and each logical drive within an extended partition is assigned a drive letter starting with C. There is a unique aspect to drive letter assignment in a FAT 16/32 environment. When assigning drive letters, primary partitions take precedence over logical drives.

For example, consider the following configuration that begins with one hard drive that has been partitioned into one primary partition and one extended partition with two logical drives.

Drive 1		
Primary	Extended	
Drive C	Drive D	Drive E

Assume that another hard drive (Drive 2) is added to the system and contains one primary partition. How would this addition affect the current configuration?

Drive 1			Drive 2
Primary	Extended		Primary
Drive C	Drive <u>E</u>	Drive <u>F</u>	Drive <u>D</u>

Note that the new drive's primary partition was assigned the drive letter "D," not "F."

Hard Drive Preparation, continued

NTFS

Windows 2000 and Windows XP can operate in a FAT32 or FAT16 environment, but most users opt for the default file system, NTFS (New Technology File System). In Windows Vista, NTFS is required to install the operating system. NTFS is more stable than the FAT system and offers such benefits as file compression and data encryption. These options are not readily available on a FAT-based system. In addition, MS Windows 2000, Windows 2003, Windows XP Professional, and Windows Vista support *dynamic volumes*, which allow partitions to be added or extended without resulting in data loss.

NTFS Partitioning and Formatting

Unlike the DOS method described previously, partitioning is not performed at the command prompt level. FDISK neither supports nor offers NTFS as a partitioning option. The process of partitioning the primary partition occurs during the installation process. This begs the question, “What if I want to add an additional hard drive to my system; how can I partition and format it with NTFS?”

The Windows Disk Management tool allows you to add, delete or modify partitions. In addition, it can be used to display general volume information including: file system type, the amount of available space and the drives total capacity. It can also be used to convert a partition from FAT16/32 to NTFS.

There are a number of ways that a partition can be formatted or reformatted to the NTFS file system. The Disk Administrator, described previously, can be used to format the partition. In addition, the partition can be formatted via My Computer. The new partition(s) display when the My Computer folder is opened. Right clicking on the partition will result in a context-sensitive menu with Format as an option. Lastly, the Format command can be executed via the command line. The command “format [drive letter] /fs:ntfs” will format an existing partition to the NTFS standard.

Hard Drive Preparation, continued

NTFS Partitioning and Formatting, continued

The following command example will convert the D: drive to NTFS.

```
convert D: /fs:ntfs
```

A reboot will be required if the desired partition is the system drive, most commonly the C: drive. Once the command is executed, the user will be prompted with a Yes/No question asking if he or she wishes to convert the specified drive to NTFS on next reboot. Selecting Yes will schedule the conversion to happen during the next reboot sequence.

Note: FAT partitions can be converted to NTFS while preserving the data; however, NTFS partitions cannot be converted to either of the FAT file systems.

NTFS File System Characteristics

NTFS utilizes the MFT (Master File Table) to track files and their associated locations on a particular volume. The MFT is similar to FAT in that it maps the location of directories and folders and is updated whenever a file is accessed, changed, deleted or added to the volume.

There are, however, significant differences. The File Allocation Table can be thought of as a static fixed-sized chart that cannot change in size. However, the MFT is much more dynamic than the FAT. The MFT is a relational database that can grow in size if necessary. The MFT is created when the drive (or volume) is formatted for the NTFS specifications.

Since the MFT has the capability of growing, a certain amount of contiguous space is reserved for MFT expansion. This area is sometimes called the “MFT Zone.” Initially, this zone is approximately 12 percent of the total volume capacity; however, the MFT can grow past that size if needed. Most NTFS volumes are no larger than 2 terabytes in size. However, the dynamic nature of the MFT allows an NTFS volume to reach 16 exabytes, which is equivalent to approximately 16,000,000 terabytes in capacity!

Hard Drive Preparation, continued

NTFS Drive Identification

Assigning drive letters in NTFS is somewhat different than in the FAT file system. Drive letters are assigned as they are discovered by the operating system. Drive letters do not change when devices are added or removed from a system.

In the following example, consider a system that consists of a hard drive with one primary partition and one logical partition.

Hard Drive 1	
Primary Partition	Logical Partition
C:	D:

Assume that a second hard drive is added to this system. This hard drive consists of a primary partition and a logical drive. The following table depicts the results of adding the new drive.

Hard Drive 1		Hard Drive 2	
Primary	Logical	Primary	Logical
C:	D:	E:	F:

Notice that the new primary partition (E:) was not assigned the drive letter D: as it would in a FAT file system.

Linux

Linux offers you the opportunity to partition a hard drive during the installation of the operating system, much like NTFS-based operating systems. Fedora, a Linux distribution based off the former Red Hat Linux, uses the Disk Druid utility to offer a graphical, mouse driven menu system that creates the partitions based on the installer's preferences.

In addition, Linux offers an FDISK program that is very similar to the DOS-based utility. It can be used to create and remove various partitions as well as view the status of a hard drive's configuration.

A partition within Linux can be formatted to a number of file systems, including:

- Second Extended File System (ext2)
- Third Extended File System (ext3)
- Reiser File System

Hard Drive Preparation, continued

File System Characteristics: Ext2

The ext2 file system offered more features than Linux's earliest file systems. These features included better space allocation, larger partition sizes (up to 16 terabytes in size), and support of long file names. Most notable, ext2 was configurable. Like most aspects of Linux, the file system could be modified to fit the needs of the end user. Ext2 was introduced in 1993 and became the standard file systems for many different versions of Linux.

Ext3

Since 2001, most Linux distributions began to use ext3 as the default file system. In comparison to ext2, ext3 provides greater reliability by taking advantage of journaling, which offers a more reliable file recovery process than ext2.

Journaling is a feature of a file system that keeps track of all items related to the main data areas of the disk. For example, when saving changes to a file, the system will inform the journal that it is about to make certain changes to a file. Once the change has been implemented, the journal entry is either marked as completed or is removed all together.

The benefit of journaling is that it can recreate anything that was lost or corrupted due to a problem or failure. For example, if a power outage occurs while you are saving a file, the journal may be referenced to complete the change after power is restored. At the very least, the original file will not be corrupted and it will be returned to its original state.

If the same scenario occurred on an ext2 system, some of the cached data may not have been written to the disk, and the file would become corrupt. After an unclean shutdown, each volume must be checked for consistency before it can be mounted on an ext2 file system.

Ext3 systems do not require a file system check after an unclean shutdown. System checks occur only in extreme circumstances such as hardware failures. Recovering from a power failure takes much less time when using an ext3 file system rather than the ext2.

Ext3 also boasts a faster read/write speed than ext2. This is because ext3 uses the journal to not only protect files, but also to optimize the hard drive's read/write head motions. As a result, the read/write process is completed in a more efficient manner.

Hard Drive Preparation, continued

Ext3, continued

Volumes using the ext2 file system can easily be converted to the ext3 standard with one command using the tune2fs program. Partitions do not have to be removed and recreated; rather the journal system is updated to the ext3 format without the loss of any data. Like ext2, the ext3 file system can manage partitions up to 16 terabytes in size.

Reiser

The Reiser file system has some similarities with the extended file systems. For instance, it can address a partition up to 16 terabytes in size. However, Reiser is considered a more robust file system. Like ext3, Reiser takes advantage of journaling for greater file reliability.

Reiser also provides the capability of online resizing. This allows the file system to automatically extend the size of the partition as needed. It only allows for extending the partition and not shrinking the partition. Although there are tools to allow both growing and shrinking of the file system while offline.

Development for Reiser version 3 has been stopped to focus development efforts to the next version of Reiser, version 4.

Ext4

The Ext4 file system is still currently under development. Some of the primary improvements are:

- Supports volume sizes of up to 1 exabyte
- Support large file sizes up to 2 terabytes
- Ability to undelete files
- Extents - A continuous area of storage reserved for a file
- Persistent file preallocation
- Online defragmentation

Hard Drive Preparation, continued

Linux Drive Identification

The drive identification process in Linux differs for the FAT or NTFS environment. Identification values are assigned based on how the device is connected to the motherboard, its jumper settings, and its relationship to the other storage devices.

File System Compatibilities

The following chart represents common operating systems and the file systems supported by each.

Operating System	FAT16	FAT32	NTFS	Reiser Ext2/3
Windows 95	Y	N	N	N
Windows 95 b,c	Y	Y	N	N
Windows 98 / 98 SE	Y	Y	N	N
Windows ME	Y	Y	N	N
Windows NT	Y	N	Y	N
Windows 2000	Y	Y	Y	N
Windows XP / Vista	Y	Y	Y	N
Linux	Y	Y	R/RW	Y

Note: Linux natively has the ability to read an NTFS file system. However, most current Linux distributions preinstall modules like NTFS-3G, which provide read-write functionality.

Key:

	= native file system
--	----------------------

This page intentionally left blank.

Lesson 2 – Floppy Drives and Removable Media

Introduction This lesson examines types of floppy disk drives and a variety of removable media storage components.

Purpose of this Lesson Understanding data storage components of a computer system is important to knowing how to safeguard information that is seized during a crime investigation.

Objectives After successfully completing this lesson, you will be able to:

- Identify the characteristics of floppy disk drive components
- Explain the characteristics of a magnetic drive
- Recognize common removable media
- Explain the characteristics of a magneto-optical drive
- Name the differences between magnetic and digital audio tapes (DAT)

In this Lesson The following table shows the contents of this lesson.

Topic	See Page
Floppy Disk Drives	D-36
Removable Media	D-38

Floppy Disk Drives

Floppy Disks

Floppy disks are storage devices used to move information from one system to another and to make file backups. There are two different types of floppy disks available:


5.25-inch Disk

This disk is a double-sided, high-density disk. They have 80 tracks with 15 sectors per track; 4 KB clusters, and either 420 KB or 1.2 MB of storage. They are nearly obsolete but may be found in older systems dated before 1994.



3.5-inch Disk

Most PCs today use 1.44 MB 3.5 inch disks for data transfer and backup, though their integration in newer computers is becoming rarer. These disks are encased in a hard plastic shell and come in three density types:



1. Double Sided/Extra High Density: 80 tracks; two sides, 36 sectors per track, 4 KB clusters; 2.88 MB of storage space.
2. Double Sided/High Density: 80 tracks; 2 sides, 18 sectors per track, 4 KB clusters; 1.44 MB of storage space.
3. Double Sided/Double Density: 80 tracks; 2 sides, 9 sectors per track, 4 KB clusters; 720 KB of storage space.

Floppy Disk Drives, continued

Floppy Drive Controllers

Today, most floppy drive controllers are now found on the Super I/O chip. In the past, the floppy disk controller was found on a separate ISA slot card. Characteristics of floppy drive controllers include the following:

- Connected to the MB for data input/output via a 34-pin ribbon cable that has a twisted segment to differentiate drive A from drive B. Unlike the controller connections for a hard drive, a floppy controller has no standard location for pin 1. Some drives locate pin 1 close to the power connector while others place it away from the power connector.
- Modern transfer rate has been standardized to approximately 1 Mb/sec.

Removable Media

Removable Media Floppy drives have become almost completely replaced by other forms of removable media that can store large amounts of data. This trend began with devices such as ZIP drives and LS 120 drives, and continues now with flash drives. Removable media can be divided into several categories: magnetic media, magneto-optical media, DVDs, CD-ROMs, USB flash drives, and external drives.

Magnetic Media Drives

Magnetic drives include the SyQuest drives and Iomega ZIP drives.


SyQuest Drives

- Uses 5.25 inch and 3.5 inch disk cartridges that use a rigid platter housed in a plastic cartridge
- Disk capacity: 5.25 inch disk holds 44 MB and 88 MB and 3.5 inch disk holds 105 MB and 270 MB
- Used with ATAPI and SCSI interface models
- Drives are not widely used but may be found in older computer systems

Iomega ZIP Drives

- Uses a proprietary disk format, which is a 3.5 inch disk cartridge with a flexible platter housed inside a plastic cartridge. These disks are about twice as thick as 3.5 inch floppy.
- Disk capacity is 100 MB, 250 MB, and 750 MB
- External models use SCSI, parallel, and USB interfaces and the internal models use SCSI or ATAPI interfaces

Iomega ZIP drive and cartridge



Removable Media, continued

Magneto-optical Drives

Magneto-optical drives provide magnetic, high density storage. Data is written to these drives by a precise laser beam that heats a tiny spot in the disk's alloy to the Curie point of 300 degrees. This action allows the molecules to be realigned when subjected to the magnetic field located on the opposite side of the disk.

To read stored data on the drive, the weaker laser beam is used to focus on the tracks. The crystals in the alloy polarize the light from the laser. This reflective beam is read by a photodiode that translates the reflective light to ones and zeros. This data is then passed on.

Super Floppy or Floptical Drive

The LS-120 drive, known as the *super floppy or floptical*, can store twenty times more data than a conventional floppy. Most super floppies hold 120 MB of data on a 3.5-inch disk and Sony has a 200 MB model. It is nearly identical in appearance to the standard floppy disks and can read/write to floppy disks. It was developed by 3M, O.R. Technologies, and Mastushita-Kotobuki Industries and uses the same ferrite materials common to floppy drives.

Floptical drives use a special optical mechanism to align the magnetic read/write heads over the data tracks. The servo track information that the laser uses to guide the read/write heads is etched into the disk by the manufacturer.

If you format a floptical disk in a standard floppy drive you will reduce storage capacity because the disk will be formatted with the standard 1.44 MB. These drives may be found in ATAPI and SCSI systems.

Removable Media, continued

Tape Drives

Tape drives are used primarily to backup entire systems. Tape drives come with their own proprietary software and newer models have ATAPI and SCSI interfaces. The two types of tape storage include magnetic tapes and digital audio tape (DAT).

Magnetic Tape Storage

- Operate much like the old reel-to-reel recording systems
- Manufactured in two sizes: 1) data cartridge (6 inches by 4 inches by 0.675 inch) 2) mini cartridges (3.2 inches by 2.4 inches by 0.4 inch) are more widely used than the data cartridges
- Stores from 2 GB to over 500 GB of data



DAT Storage

- Records in a digital format using magnetic technology
- Tape size is a little smaller than a cassette tape
- Available in two primary formats: 1) 4mm cartridge that stores up to 36 GB of uncompressed data and 2) 8mm cartridge that stores up to 80 GB of uncompressed data
- A third format, Digital Linear Tape (DLT), stores between 35 GB and 800 GB of data and is only used with high-end data servers

Removable Media, continued

Compact Disc Read Only Memory (CD- ROM)

CD-ROM is a disc format used to hold text, graphics and hi-fi stereo sound. It is very similar to an audio CD, but uses a different format for recording data. An audio CD player cannot play CD-ROMs, but CD-ROM players can play audio discs.

How Data is Stored on CD-ROM

Information is stamped into the disc when manufactured as a series of pits and lands. The pits are the small raised areas in the tracks on the CD-ROM separated by flat spaces (lands). The back of the CD-ROM is a highly reflective surface. When the laser beam shines up through the disc, the light is reflected back. When the light hits the edge of a pit, the light is not reflected back. When the light hits a land, the light is reflected back. The microprocessors in the drive translate the light and dark transitions as 1s and 0s.

Compact Disc Recordable (CD-R)

CD-R is a form of CD-ROM technology that writes to a disc. Each disc can be written to only once. A CD-R drive cannot stamp pits into the disc. Instead, it uses a layer of dye to simulate the pits and lands. To record, the laser burns the dye to simulate pits and does not burn areas for lands. To read, the laser beam is reflected back from the non-burnt areas (lands) and is not reflected back from the burnt areas (pits). The drive interprets both.

Compact Disc Re-writable (CD-RW)

Unlike a CD-R that can only be written to once, a CD-RW can record, erase, and re-record data on the same disc up to 1,000 times. Instead of stamped pits or dye, CD-RWs use a phase change technology that involves a crystalline layer. The process uses a high intensity laser pulse that turns the crystalline natural state (reflective) to an amorphous one (dull). When the disc is read, it uses this dull/reflective state to simulate the pits and lands. To erase and re-record data, CD-RW uses a medium intensity pulse to restore the crystalline surface to a reflective state.

Removable Media, continued

DVD Introduction DVD is commonly termed Digital Video Disc as a means to describe the technology. However, a DVD is far more capable than just storing and playing digital video.

DVD is a group of disk formats that store data, video, or audio information with capacities between 4.7 GB and 8.54 GB. A DVD is generally a disk that is 120 mm in diameter and 1.2 mm thick. Reflective surfaces embedded in polycarbonate resin hold data as pits and lands in tracks that are much smaller and closer together than those on a CD. All DVD drives feature backward compatibility to read CD media.

DVD-ROM Data, be it audio, video, or software, is molded into the disk when manufactured. Unlike CD media, DVD media has its reflective data layer more deeply protected in the structure of the disk.

User-Recordable DVD Data is written or “burned” into a disk by the user with a write capable drive. Currently, there are six different formats of writable DVD. You must use compatible media in the appropriate drive to make a recording. Once written, all DVD disk types (except DVD-RAM) can be read in other DVD drive types.

Removable Media, continued

Recordable DVD Media Types

DVD-RAM	Once formatted, this disk together with its drive can be treated as if it were a hard disk drive. Re-writable up to 100,000 times. For more capacity, use another blank disk.
DVD-R	Write-once, read many times
DVD-RW	Re-writable up to 1,000 times
DVD+R	Write-once, read many times
DVD+RW	Re-writable up to 1,000 times
DVD+R DL	Dual layer write once, read many times

Note: Drive types are typically grouped as DVD-R/RW/RAM, or DVD+R/RW/DL.

DVD Capacity Comparison

The table below lists a comparison of DVD capacities.

Format	Capacity	Type
Single Sided, Single Layer	4.7 billion bytes	“DVD-5”
Single Sided, Dual Layer	8.5 billion bytes	“DVD-9”
Double Sided, Single Layer	9.4 billion bytes	“DVD-10”
Double Sided, Dual Layer	17 billion bytes	“DVD-18”

HD-DVD and Blu-ray

Introduced in 2006, HD-DVD and Blu-ray are two rival High Density DVD formats. Each format is backed by competing groups of equipment manufacturers.

HD-DVD has a data capacity of 15 GB for single-layer disks and 30 GB for dual-layer disks. In late 2007 a new format was approved, allowing for three layers to store a capacity of 51 GB.

Blu-ray disks can hold 25 GB for single layer disks and 50 GB for the dual-layer disks. Laboratory tests have successfully tried up to eight layers for a total of 200 GB per disk.

Removable Media, continued

USB Flash Drives Flash drives are mass storage devices that allow you to transport files between devices. Flash drives offer a very compact, non-volatile storage option for data. The transfer rate of data is dependent on the type and speed of interface used for the device.

USB Flash Drives

- Small, easily transported, and rewritable
- Uses USB 1.0, 1.1, or 2.0
- Stores up to 64 gigabytes

Memory Cards Memory cards are flash memory storage devices that are common in many electronic devices. Some devices that may support memory cards are: digital cameras, cell phones, portable audio devices, game consoles, mobile computers, etc. The large storage capacities and small footprints make this storage technology an excellent companion to many portable electronic devices.

Compact Flash - Introduced by Sandisk in 1994. With storage capacities up to 64 GB, these are commonly found in digital cameras or mini hard drives. Available in two variations: Type I and Type II.

Secure Digital (SD) - Developed by Sandisk, Matsushita, and Toshiba. Secure Digital format supports up to 16 GB of storage and up to 32 GB with high-capacity SDHC cards.

MiniSD - Provides storage capacities up to 4 GB. These devices can be used in a typical SD slot with the use of an adapter. High-capacity MiniSDHC cards can provide for up to 32 GB of space.

MicroSD - These cards are about the size of a penny and provide storage capacities up to 8 GB. These cards can be used in the larger SD and MiniSD slots with the use of an adapter. A high capacity implementation, microSDHC, allows for capacities up to 32 GB.

Memory Stick - This memory is developed by Sony. The term “Memory Stick” can refer to the whole memory stick lineup of products. These include the Memory Stick Pro, memory Stick Duo/Pro Duo, Memory Stick Micro, and Memory Stick Pro-HG.

Removable Media, continued

External Drives

External drives are available in USB, FireWire, and eSATA. These drives are hot swappable and easily transported from one system to another.

USB and FireWire connections are recognized by operating systems such as Windows XP and many versions of Linux. The eSATA has a special adapter and cable, which eliminates the protocol issues of converting the drive's native signal to the serial port adapter. This allows higher transfer rates, but restricts the portability of the drive to computers equipped for eSATA.

This page intentionally left blank.

Appendix E

Input/Output Components

Overview	This module examines the various components involved in the transfer of data into and out of a computer system.
Purpose of this Module	The purpose of this lesson is to introduce the data transfer components of computer systems. Student will also be taught how the various components interact in order to move data.
Objectives	<p>After successfully completing this module, you will be able to:</p> <ul style="list-style-type: none">• Recognize basic input devices such as the keyboard, mouse, scanner, and modem• Explain how monitors and video display adapters work• Identify the various input/output ports found on a PC• Define interrupts, IRQs, direct memory access, and device drivers• Recognize SCSI devices and connectors
In this Module	The following table shows the contents of this module.

Topic	See Page
Lesson 1 – Input/Output Devices and Ports	E-3
Lesson 2 – BIOS and System Initialization	E-23

This page intentionally left blank.

Lesson 1 – Input/Output Devices and Ports

Introduction This lesson introduces the basic input/output devices including the keyboard, mouse, scanner, monitor, printer, and modem.

Purpose of this Lesson You will gain a broader understanding of common input/output components and how they work.

Objectives After successfully completing this lesson, you will be able to:

- Recognize input devices such as the keyboard, mouse, scanner, and modem
- Explain how monitors work and be familiar with video display adapters
- Identify the various input/output ports

In this Lesson The following table shows the contents of this lesson.

Topic	See Page
Input/Output Overview	E-4
Input Devices	E-5
Output Devices	E-8
Input/Output Ports	E-13
Modems	E-19
PC Cards	E-21

Input/Output Overview

Overview

Most of the hardware components found on a PC can be categorized as either an input or an output device. Modems are the exception because they transmit and receive information. Devices included in each category are:

Input Devices

- Keyboard
- Mouse and other pointing devices
- Game controllers
- Scanners

Output Devices

- Monitors
- Video Adapter Cards
- Printers

Input Devices

Input Devices A computer system includes several data input devices. Those defined here include the keyboard, pointing devices, and scanner.

Keyboard The keyboard is the primary input component for using a computer. Many new features have been added to the keyboard in recent years such as additional short cut keys. However, the basic way a keyboard works remains the same since the first IBM PC in the early 1980s.

There are many styles of keyboards with various built-in functions including pointing devices such as trackballs and special keys that are programmed to launch programs or Web sites.

How Keyboards Work

In brief, the keyboard works using a grid of embedded circuits called the key matrix. When a key is pressed, the change in current in the circuit that is associated with the key is picked up by the keyboard's microprocessor. The microprocessor generates a scan code that is read by the computer's BIOS. The BIOS translates the scan code into ASCII code, which is then recognized by the application software, such as word processing software.

Pointing Devices Computer pointing devices come in many shapes and sizes. All are used to control a computer by pointing to images instead of using the keyboard to type commands. The mouse was the first device developed by Xerox in the early 1980s and remains the standard. Other current devices include the touchpad, trackball, trackpoint, and game controller.

Mouse

The mouse is the standard pointing device used to navigate your system's graphical user interface (GUI). In short, the internal electronics of the mouse sends signals through the attached cable to the computer's operating system. The operating system then translates the signal to move the onscreen cursor.

Input Devices, continued

Pointing Devices, continued

Touchpad

The touchpad is a fixed pointing device that uses the movement of a finger on the touch-sensitive pad to control the location of the cursor in the GUI. It replaces a standard mouse in laptop and notebook computers where space is limited. It is also known as the glidepoint.

Trackball

The trackball is basically an inverted mouse with the ball located on the outside of the housing. It contains the same connectors as found in a standard mouse.

Trackpoint

The trackpoint, also called a pointing stick, is a small rubber component embedded in the keyboard between the G and H keys on a laptop computer. To operate it, you move the stick to move the cursor.

Gyroscopic Mice

Gyroscopic mice feature gyroscopes to detect movement while the mouse is moved in mid-air. Such mice are normally shaped similar to remote controls and allow for wireless control of a computer from long distances. They are commonly used for conference room and home theater computers.

Game Controllers

Game controllers, also called joysticks, are available in a wide variety of designs such as steering wheels, flight simulator controls, and pressure pads. All are input devices for game software. The joystick has an upright stick that controls movement by providing the X-Y axis coordinates of the stick.

Joysticks come in digital or analog compatible formats. They connect through game ports or the USB port. The more expensive controllers can be programmable with macros to allow specific series of actions or commands to be transmitted with a single button.

Input Devices, continued

Scanner

Scanners are used to translate text or pictures into a digitized form the computer can read and store. All scans are delivered to the computer as a digital image. Scanned pictures are code that can be used by a graphics software program. Printed type can be converted into editable text when scanned using optical character recognition (OCR) software.

There are three main types of scanners:

1. Hand-held scanners: Offer low resolution, 200-400 dots-per-inch (DPI), with 8-bit imaging. The scan window is only five to seven inches wide. Therefore, wide images require image resectioning.
2. Flatbed scanners: Provide photo quality resolution. Typed sheets and pictures are placed facedown on the large scan window. Advantages include high resolution and the ability to scan thick documents such as books. These scanners usually come with their own proprietary image processing software.
3. Sheet-fed scanners: Scan one standard-size sheet at a time. Resolution can be the same as flatbed scanners.

Output Devices

Output Devices A computer system includes several data output devices. The following describes output devices including monitors, video display adapters, and printers.

Monitors The main video display components of a PC include the monitor and the video adapter, also known as a video card. A monitor refers to the display screen. Monitors are classified by the following criteria:

- Diagonal screen size measured in inches. Desktop PC monitors range in size from 14 inches to over 40 inches and LCDs in laptops generally range from 6 inches to 21 inches. Monitors are usually evaluated on two prime factors: size and resolution.
- Display resolution measured in pixels
- Refresh rate ranging from 60 Hz to 100 Hz

Monitor Technologies

The three types of monitor technologies are:

- Cathode ray tube (CRT) monitor: Most common type of monitor for desktop systems until the early 2000s.
- Liquid crystal display (LCD) monitor: Used in laptops as well as high-end desktop monitors.
- Organic Light Emitting Diode (OLED) monitor: Used in laptops as well as desktop monitors.

Output Devices, continued

Benefits of CRT Monitors


While the sales of CRT monitors have decreased dramatically in recent years, they still hold a few major advantages over LCD monitors. The color display on CRT monitors is dramatically better than LCD, providing a greater color range and depth, which is beneficial to graphic artists. LCDs also suffer from blurring during fast motion, and slower response times in redrawing the screen. These two issues make high-performance CRT monitors beneficial to avid gamers.

CRT Monitor Resolution

Resolution of a monitor is measured by pixel density. The more pixels it has, the clearer the image. Pixel density is expressed in dots per inch (DPI). Other notable resolution characteristics of monitors include the following:

- **Dot Pitch:** The dot pitch is the distance in millimeters between the same color phosphor dots/squares on the screen. The less the distance between dots, the higher the screen resolution. Likewise, the smaller the dot pitch, the greater the clarity. For example, an average resolution is approximately .28 mm, whereas a high resolution is .26 mm or less.
- **Refresh Rate:** Speed at which a monitor completes drawing the screen. The speed is measured in Hz. The higher the refresh rate, the sharper the image will be with less screen flicker. The minimum PC refresh rate is 60 times per second.

Example of Dot Pitch



Output Devices, continued

Monitors: Liquid Crystal Display and Organic Light Emitting Diode

LCD monitors are found in calculators, handheld devices, televisions, PCs, and notebook computers. LCDs produce different colors on the screen by using a combination of filters and electrically charged liquid crystal cells to filter and change the angle of the light passing through the panel. Transistors electrically charge the crystals. The number of crystals used is proportional to the number of pixels the screen can display.

LCDs generally use less power than their CRT counterparts. However, they are more expensive to manufacture than CRT's and do not have the same clarity, contrast ratio, or response time. There are two categories: passive matrix and active matrix LCDs.

Passive Matrix LCDs

A passive matrix screen has transistors along the edge of the screen that are connected by conductors. The intersection of conductors forms a pixel. If the screen has 1,024 x 768 pixels, there are 1,024 transistors along the horizontal edge and 768 transistors along the vertical edge to electrically charge the crystals. By placing the transistors along the edge of the screen, passive matrix LCD monitors are less bright and cheaper to manufacture than active matrix monitors.

Active Matrix LCDs

Active matrix LCD monitors use the same system of filters and liquid crystal to create the image as passive matrix LCDs, but have a transistor for every pixel. As a result, active matrix LCDs require a greater number of transistors. For example, a resolution of 1,024 x 768 requires 1,892 transistors for a passive matrix monitor and 786,432 transistors for an active matrix monitor. The active matrix monitors provide greater clarity and brightness.

Organic Light Emitting Diode (OLED) Monitors

OLED monitors are a newer evolution of LCDs. Unlike LCDs, they do not require a backlight to display pixels, therefore greatly reducing power requirements. Organic LEDs have an organic substrate that glows when you put an electric current behind it. The material used is also quite flexible, allowing for its use in unique applications, such as clothing.

Output Devices, continued

Blurring the Line Between Monitors and Televisions

As consumer televisions have evolved in recent years, the distinction between TVs and monitors has greatly eroded. Many modern televisions, whether LCD or Plasma, now feature VGA and DVI inputs to allow direct connections from computers. It is possible to connect a computer directly to even a 65" LCD display for regular computing.

MultiTouch Monitors

MultiTouch monitors are an improvement over regular touch screen monitors by allowing multiple fingers or styluses to interface with the screen at a time. An example of this technology is Microsoft's Surface Computer.

Video Display Adapters

Video display adapters, also called video cards, are either expansion cards or chipsets in the motherboard that display images on the monitor. Video cards can include the following components:

- Independent BIOS and chipset to control image on the screen by first writing the data to the video RAM
- Video RAM with capacities even as high as 1 GB
- Video processor (does not depend on the CPU to process signals)
- Optional TV-In/Out connector that redirects the display to a standard television set or video processing/capturing equipment

Example of Video Card



Output Devices, continued

Printers

A printer is a device that prints illustrations, charts, or text on paper. The most prevalent types of printers are:


- **Dot Matrix:** A high-speed printer that uses a grouping of mechanical pins in the print head. The text is displayed in rows of dots.
- **Ink Jet:** A higher resolution type printer when compared to the dot matrix model. Printing occurs when tiny nozzles spray magnetically directed, ionized ink on the paper.
- **Laser printer:** The highest resolution printer. Print technology uses a toner set and laser tracing to transfer the image to paper.

Input/Output Ports

Recognizing I/O Ports

You should be able to recognize the various input/output ports on a PC. Knowing their form and location will ensure that you make the proper connections.

Computer Ports



Mouse and Keyboard Ports

The characteristics of the mouse and keyboard ports are:

- A PS/2 mouse and the keyboard use a barrel shaped 6-pin mini port
- The two ports look exactly alike unless they are color coded or labeled
- A mouse could require a 9-pin serial or USB port
- Older keyboards connect using a 5-pin barrel-shaped connector that is slightly larger than a PS/2 port. Connectors are found on AT or older motherboards.

PS2 Mouse or Keyboard Port



AT style 5 pin
keyboard connector



PS/2 mouse or
keyboard connector


Input/Output Ports, continued

Serial Ports (Com 1 & Com 2)

A serial port sends only one bit of data at a time over one wire, and is typically used to connect older devices to a computer, such as modems, touch screens, and proprietary devices. PCs have a maximum of four serial ports that are referred to as COM 1, COM 2, COM 3, and COM 4. These ports are divided into pairs: COM 1 is paired with COM 3 and COM 2 is paired with COM 4. Each pair shares the same system resources.

Because they share resources, the ports that comprise a pair cannot be used at the same time. For example, COM 1 and COM 3 cannot run at the same time.

Serial ports have two rows of either 9-pins or 25-pins.




Parallel Port

Most PCs have one parallel port. It is mainly used to connect a printer although it can be used for certain external drives such as Zip or Jazz drives. The parallel port is much faster than the serial port because it can send one byte of data at a time, instead of just a single bit.

The port is either a 25-pin parallel port or a 36-pin Centronics port. The SCSI-1 parallel port uses a 50-pin Centronics port. You should not confuse the different types of parallel ports.

Example of Parallel Port



Input/Output Ports, continued

USB Port

Universal serial bus (USB) port enables many different components to be connected to the computer sharing the same port using a connection configuration called daisy chain. The USB port is a small, flat, rectangular connection that requires a four-wire cable. The ports are usually found in pairs.

The most important feature of the USB is its ability to hot swap or change devices without turning off the computer. In theory, 127 devices can be daisy chained.

The transfer rate of USB 1.1 is 12 Megabit/sec (1.5 MB/sec) and the transfer rate of USB 2.0 is 480 Megabit/sec (60 MB/sec).

The symbol signifying the USB port can be found on the back of the computer.

USB Symbol



FireWire and i.LINK Ports


FireWire made by Apple Computers and i.LINK made by Sony are the manufacturer brand names for IEEE 1394 ports. These ports are an alternative to a USB port and transfer data at a much higher speed than USB 1.1, generally 50 MB/sec. They have either 4-pin or 6-pin connectors with 6-pin port considered the standard. These ports allow up to 63 devices to be daisy-chained to one connector.

Input/Output Ports, continued

Game Controller Port

Game controller port connects a joystick or other type of game controllers to the PC. The port has 15 pins distributed in two rows.


Game Controller Port



VGA Video Display Ports

One of the most common video connector ports is the VGA, which transmits an analog video signal. It has 15 pins distributed in three rows, forming an arrow that points in one direction.


VGA Video Port



DVI Video Display Ports

Evolving from the current VGA standard is DVI, or Digital Visual Interface. Unlike VGA's analog signal, DVI provides a pure digital signal for higher quality displays. However, DVI has been released in a number of different formats, few of which are cross-compatible.

DVI-D: This format is true digital video. The connection is used from digital video cards to a digital display. This produces higher quality output since the signal isn't being converted to another format.




Input/Output Ports, continued

DVI Video Display Ports, continued

DVI-I: These cables can be used to carry either digital or analog VGA signals. A device with a female DVI-I connector can accept DVI-D cables along with DVI-I cables.




Along with the different DVI connectors, each is also available as either Single-Link or Dual-Link. Single-link is the most common form, but larger displays, such as 30" LCD monitors, require Dual-Link. Shown below is an example of a DVI-I Dual-Link connector.



HDMI

High Definition Multimedia Interface (HDMI) is an interface used to transmit video, audio, and/or device controlling signals. HDMI supports 1080p high-definition signals and up to 8 channels of uncompressed audio. The 8 audio channels allow for the use of 7.1 surround sound.


HDMI uses transition minimized differential signaling (TDMS) to encode the signal to help protect it from degradation during transmission. The source of the signal encodes the digital signal to minimize the amount of data required to be sent over the cable. One of the twisted pairs will send the encoded signal and another pair will send the inverse of the signal. The receiving device will decode the signal then compare original and inverse streams. The device will use that comparative information to compensate for any degradation of signal during transmission.



Input/Output Ports, continued

DisplayPort

DisplayPort, supporting 2560x1600 resolutions, is a license and royalty free competitor of HDMI. A device with the DisplayPort will be able to connect to any existing DVI, HDMI, VGA, or CRT monitors.



Sound Ports

Sound cards and sound chips usually have four round 1/8 inch mini TRS jack ports. These four ports are color-coded and are defined as line-in, line-out, microphone, and speaker-out. Other sound ports may be found on the PC including the following ports:

- RCA stereo port carries sound to stereo speakers that are not specifically made for the PC
- MIDI port allows the connection of a musical instrument such as a synthesizer to the PC
- Digital audio output

Modems

Modems

Modems handle communication that is transmitted over telephone lines between computer systems. Most modern modems have fax capabilities.

PCs are digital devices and the telephone system is an analog device. The modem is the component that converts (modulates) the PC's digital code to analog so it can be sent over phone cables. Likewise when receiving information, the modem converts (or demodulates) analog signals to digital code before transmitting data to the PC. The transmission mode takes two separate forms:

- **Asynchronous Mode:** This mode of transmission uses transmission stops, which are nothing more than the periodic cessation of data transfers. The data is sent in a start and stop bit fashion.
- **Synchronous Mode:** This mode of transmission is a constant data stream.

Modems also use two methods to facilitate the transfer of information:

- **Data Compression** – Speeds data transfer
- **Error Correction** – Allows the modem to detect errors in transit and either correct these errors or resend the data

Modems have standardized speeds and protocols. Speed is measured in baud or bits per second and varies depending on the model. The standard transmission protocol for modern modems is V.92 protocol that transmits at a speed of 56K.

Types of Modems

External Modems

- More expensive than internal modems
- Include a separate chassis and power supply
- Readily identifiable by the LEDs on the front panel
- Connect to the serial or USB port. This factor may limit the modem from reaching its full transmission capabilities.


Modems, continued

Types of Modems, continued

Internal Modems

- Internal modems are contained on a PCI or ISA card
- An internal modem will contain separate communications circuitry and is not limited by the serial port connection in any way
- Internal modems are less expensive than external modems

Example of Internal Modem



Caution: You should never attempt to connect a real analog modem to a multi-line or digital connection. Voltage requirements are different and a connection will destroy the modem.

Winmodems

Winmodems are internal modems that are configured by device drivers in an operating system rather than by physical jumpers. They use the CPU to process data, instead of an onboard chip, which can degrade system performance. As they require system drivers to operate, these modems usually only work with Microsoft Windows.

PC Cards

PC Cards

PC cards, also known as *PCMCIA* (Personal Computer Memory Card International Association) cards, are metal-encased expansion cards about the size of a credit card. These cards, most often used in laptop computers, add various functions to the system including:

- Fax and modem
- Mini-hard drives
- Network Interface Card (NIC)
- Small Computer System Interface (SCSI) host adapters
- Additional RAM and ROM

PC cards are manufactured in the following four types and all cards have the same dimensions (3.4 inches by 2.1 inches). The differences are in card thickness and functions. Different card types require different card ports, although Type I cards can be read by a Type II reader and Type II cards may be read by a Type III reader.

- Type I is primarily used to add more memory (3.3 mm thick)
- Type II is used for modems and NICs (5 mm thick)
- Type III is used to add a mini-hard drive (10.5 mm thick)
- Type IV is used to add a high-capacity drive. Type IV is not officially recognized by the PCMCIA and there are no official standards for implementation.

This page intentionally left blank.

Lesson 2 – BIOS and System Initialization

Introduction Earlier you learned that the BIOS is part of the motherboard and its basic function. You will now learn how it interacts with all parts of the computer and launches the operating system.

Purpose of this Lesson You will understand how the motherboard uses the BIOS to initialize the hardware and start the operating system.

Objectives After successfully completing this lesson, you will be able to:

- Define the role of the BIOS
- Discuss what POST codes are
- Explain BIOS and the concept of Plug and Play
- Explain how a system boots

In this Lesson The following table shows the contents of this lesson.

Topic	See Page
Motherboard Components	E-24
BIOS Information	E-27
The Boot Process	E-30
The Master Boot Record	E-32

Motherboard Components

Motherboard BIOS

The motherboard *basic input/output system (BIOS)*, also called the system BIOS, is considered to be the heart of the computer because it controls communications between computer hardware and the operating system. The system BIOS is also referred to as ROM BIOS because the code is contained in a non-volatile, read-only memory (ROM) chip.

The system BIOS contains software called firmware. *Firmware* provides the basic input/output instructions to boot the computer and handles several important functions including identifying hardware currently installed in the PC, determining which device will boot the PC, and installing basic drivers for keyboard, video, and disk drives prior to the operating system loading.

CMOS

The CMOS chip, or complementary metal-oxide-semiconductor, refers to a particular non-volatile memory located on all computer motherboards that is used to store information from the BIOS. While modern computers don't use a technical CMOS chip anymore, the term is still used to refer to any memory on the motherboard used to store motherboard settings. The CMOS is used to store all values that are configured within the system BIOS including: hard drive geometries, date and time, and startup settings.

To maintain the storage of information, even while the machine is powered off, many CMOS chips are powered by an external battery. In older machines a typical 9-volt battery is common, but in modern computers a lithium watch battery has become the new standard. Such batteries tend to last for up to 10 years.

Functions of Motherboard BIOS

The four major functions of BIOS are:

Function 1: Contains and executes the POST

Function 2: Contains the CMOS setup program

Function 3: Executes the bootstrap loader

Function 4: Contains and runs the initial boot device drivers

Motherboard Components, continued

Functions of the Motherboard BIOS, continued

Function 1: Contains and executes the power-on-self-test (POST.) This is a generalized, hard-coded diagnostic utility that examines crucial components of the system prior to the boot procedure. This program tests the RAM, video card, keyboard, and disk drives. If POST finds a fatal error that prevents the system from booting, POST sends the following alerts:

- A series of audio codes or beeps prior to the initialization of video.
- On-screen text messages after initialization of video.
- Numeric codes sent to an internal I/O port address. These codes can only be read by a special PCI or ISA card.

When you encounter these warnings, you can either continue to boot, if possible, or enter the setup to reconfigure the system and resolve the error. In addition, you can view the codes on the BIOS manufacturer's Web site.

System BIOS maintains information about all the legacy and Plug and Play devices that are discovered via POST.

Function 2: Contains the CMOS setup program. When a computer boots, you can generally find instructions on-screen such as "Press F1 for Setup." Setup is a program that is run by the motherboard BIOS and the data generated by Setup are stored in a separate chip called CMOS. CMOS stores data such as date and time, which device is used to boot the PC, and the ability to enable or disable many of the motherboard functions.

Function 3: Executes the bootstrap loader. This is a code that queries CMOS to determine the boot device, locates the Master Boot Record for that device, locates the operating system, boots the system, and then gives control of the computer to the OS. In short, this function performs a quick check of the computer system to ensure that it is ready to accept the operating system.

Motherboard Components, continued

Functions of Motherboard BIOS, continued

Function 4: Contains and runs the initial device drivers needed to boot the system. A device driver is a small piece of software code that allows the operating system to communicate with a physical device. The device drivers loaded by the system BIOS are very basic drivers for keyboard, video, and disk drives. In fact, the system BIOS contains a table of data for several types of disk drives.

The BIOS interrogates the drive and uses the information in the table to identify the drive and to facilitate access to the drive. If the settings are not found through this identification process, the drive configuration settings can be entered manually through the Setup program.

BIOS Information

BIOS Setup

All X86 motherboard architecture personal computers have BIOS. The BIOS contains a program called Setup that allows users to configure low-level system settings of the computer. Typing a certain combination of keystrokes during the boot process accesses Setup. These keystrokes differ among the different BIOS manufacturers. When initial program load (IPL) ends (immediately before the computer loads the OS), the BIOS will generally pause for one to two seconds and display the keystroke combination that initializes Setup.

For most computers, you press the Delete key to access Setup. Others (especially proprietary systems such as Compaq and Packard-Bell) may require pressing these keys: F1, F2, F10, CTRL-A, or CTRL-S. If you cannot determine the keystroke to activate Setup, refer to written documentation prepared by the BIOS manufacturer. BIOS settings configured by Setup are actually stored in the CMOS chip as standard practice for most systems.

BIOS Menu System

Setup usually contains a menu system that allows easy access to system settings. Data and options are divided into categories of relevance. While searching for data, it may be necessary to search through the various menus. However, be very careful of what keystrokes you use. Keystrokes to change settings and navigate the menu differ, depending on the manufacturer. For example, the arrow keys that may navigate in one Setup program may actually change configuration values in another.

When exiting the setup program, be sure to select the exit option that will not save any changes, so as to not overwrite the original settings.

BIOS Information, continued

In Cases of Password Protection

A password protected BIOS is a security measure intended to prevent trespassers from viewing vital details of the computer's configuration. Some motherboards have a built-in backdoor password that can be used to gain access to the BIOS at any time. However, for many motherboards, the only way to quickly circumvent this security is to completely erase the CMOS (which, in turn, erases the suspect's BIOS settings.) Here are common ways to erase the CMOS:

1. Locate and change the CMOS "clear jumper" on the MB
2. Remove and replace the CMOS battery

Note: If you are forced to erase the BIOS settings, record this in your notes. Skip the steps of recording BIOS information.

Vital Information to Record Time and Date

To create a case profile, you need to gather as much pertinent information as possible from the suspect's computer. Primary data includes the date and time of the computer. Document this data and also note any discrepancies between the values on the computer and the actual time and date.

Boot Order

This sets the order in which bootable devices are scanned by the BIOS. This may be changed in Setup very easily, and the order is generally designated by "First," "Second," or "Third." Investigators should document the original order as set by the suspect.

If the "First" boot device is the network card, this may be an indication that the alleged crimes did not occur on this hard drive, but on a network server. Record this fact in your notes.

BIOS Information, continued

Hard Drive Information

You need to record the configuration of the hard drives and storage drives in the BIOS. Most computers have the capability to control four IDE devices at once or two per connector. Document settings for each device. Possible settings are “Auto-Detect,” “User-Defined,” or “None.”

If set for “User-Defined,” record all the numerical settings that are entered, including the Cylinders, Heads, Precomp, and whether or not LBA is in use. Cross-reference this information with that which was obtained during the internal examination (e.g., information from hard drive labels, etc.) and note any discrepancies.

Note: This information will be important if you want to restore the hard drive image.

The Boot Process

Boot Process Overview

All computers are designed to start in a predictable way from the moment you press the Power On button until the moment the operating system loads. The list below outlines the steps.

Boot Process	
Activity	Description
1. Power on	When you press the Power On button and the power supply reaches working voltage the Power_Good line is set to true on the microprocessor.
2. CPU looks at ROM for basic instructions (BIOS)	When CPU receives power good signal it starts to load the program at memory location 0000h, which is the start of the BIOS.
3. System BIOS loads	
4. BIOS initiates POST	POST = Power On Self Test
5. POST checks RAM and then Video. If either of these have a problem, the result are various beep codes from the speaker. <i>From this point forward, errors are reported with text messages displayed on the monitor.</i>	Depending on the manufacturer, POST error codes vary. The following web site has many of the industry error codes outlined for you. www.bioscentral.com
6. When RAM and Video pass the POST test, a single beep occurs. The single beep exists simply to indicate that the diagnostic speaker is working. A malfunctioning speaker will prevent audible beep codes.	You will begin to see text on the screen. The rapid numbers flashing indicate an in-depth RAM check. The screen will indicate the BIOS manufacturer and version number.

The Boot Process, continued

Boot Process Overview, continued

Boot Process, continued	
Activity	Description
7. POST then checks keyboard.	If an error occurs, a text message generally displays the on-screen
8. The BIOS checks memory location 0472h for the value 1234h which indicates a 'warm start', or that the reset button was pressed.	If this is a warm start then the BIOS shortens the check and load process to only those things necessary to start the system, versus a full diagnostic check.
9. CMOS data is compared against new current configuration data. Drives spin, lights flash, and sounds are heard as circuits are tested and prepared for system start-up.	If there is a problem with the CMOS battery, you will get a text message.
10. Finding no major hardware errors, BIOS turns the process over to the operating system boot loader on the default boot device.	The boot loader starts and the search for the Master Boot Record begins.
11. The boot loader learns the boot sequence from (e.g. A: C: CD-ROM, etc.) and looks for the Master Boot Record on that device. For hard disks, the boot loader looks for a partition table. The partition table will have a pointer to the MBR on the primary, active partition	
12. The MBR contains the first file needed to start the operating system (IO.SYS in Windows 9x, boot.ini in NT).	
13. The whole process is turned over to the OS and you see splash screens, etc.	

The Master Boot Record

The MBR Defined Since the early days of the disk operating system on personal computers the last thing the boot loader does is search for the Master Boot Record (MBR) for the operating system loaded on the computer. Originally this was found at cylinder 0, head 0, sector 1 of the default boot drive.

Historically this was usually the floppy disk drive.

Today just about any storage device can be configured in the BIOS as the default boot device, including USB drives, CD's, DVD's and even booting from the network interface.