



Education

First On the Digital Scene: A Forensic Primer for Storage Professionals

Richard Austin CISSP

- The material contained in this tutorial is copyrighted by the SNIA.
 - Member companies and individuals may use this material in presentations and literature under the following conditions:
 - ◆ Any slide or slides used must be reproduced without modification
 - ◆ The SNIA must be acknowledged as source of any material used in the body of any document containing material from these presentations.
 - This presentation is a project of the SNIA Education Committee.
 - Neither the Author nor the Presenter is an attorney and nothing in this presentation is intended to be nor should be construed as legal advice or opinion. If you need legal advice or legal opinion please contact an attorney.
 - The information presented herein represents the Author's personal opinion and current understanding of the issues involved. The Author, the Presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.
- NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.**

➤ **First on the Digital Scene**

- ◆ With most enterprise information concentrated within storage networks, the likelihood of a storage administrator being "first on the digital scene" of an intrusion, a crime, a policy violation or an e-discovery request is becoming almost a certainty. This presentation presents a whirlwind tour of the requirements, processes and procedures for collecting and preserving digital evidence.

•**First on the Digital Scene**

With most enterprise information concentrated within storage networks, the likelihood of a storage administrator being "first on the digital scene" of an intrusion, a crime, a policy violation or an e-discovery request is becoming almost a certainty. This presentation presents a whirlwind tour of the requirements, processes and procedures for collecting and preserving digital evidence.

- I've seen CSI but what is this “forensics” thing?
- A Primer on Evidence
- Policy – the first step
- The Forensics Process
- Collecting Evidence
 - ◆ Data on Disk
 - ◆ Log Records

Why does it matter?

➤ Regulatory and other legal obligations

- ◆ New Federal Rules of Civil Practice (FRCP) specifically recognize the importance of Electronically Stored Information (ESI)
- ◆ Support a dismissal decision and defend against a wrongful termination action
- ◆ Support or defend against other legal actions
 - Theft of intellectual property
 - Tampering, willful-destruction, etc

➤ Criminal Investigation

- ◆ Fraud, terrorism, etc




What is “forensics?”

- Webster – from L *forensis* fr *forum*. Belonging to, used in or suitable to courts of judicature or to public discussion and debate
- Saferstein – Forensic science in its broadest definition is the application of science to law
- Computer forensics – using accepted methods and procedures to properly seize, safeguard and analyze data. (Kroll Ontrack)

- An item does not officially become a piece of evidence until a court admits it as such
 - ◆ Opposing counsel can challenge this admission
- The word “evidence” is used here as a shortcut for “item of potential evidentiary value”
- Forensics practice is concerned with identifying, collecting and analyzing these items without compromising their potential to be admitted as evidence in a court of law



A Serious Business

- Evidence may save or cost an enterprise millions of dollars 
- Evidence may deprive a person of their liberty in a criminal matter (which now includes corporate governance) 
- Evidence may deprive a person of their livelihood if it leads to termination of employment 

Because of the serious nature of these consequences, the legal system imposes stringent requirements for information to be used as evidence in their proceedings.

- Relevant
 - ◆ Has an important role in deciding a question of fact
- Authentic
 - ◆ The “real” thing
- Integrity Preserved
 - ◆ From collection through analysis to presentation
 - ◆ Unbroken chain of custody



This is a private system

Policies that deal with
the expectation of
privacy and continuous
monitoring

I have read and understood

Policies communicated and
agreed to in writing

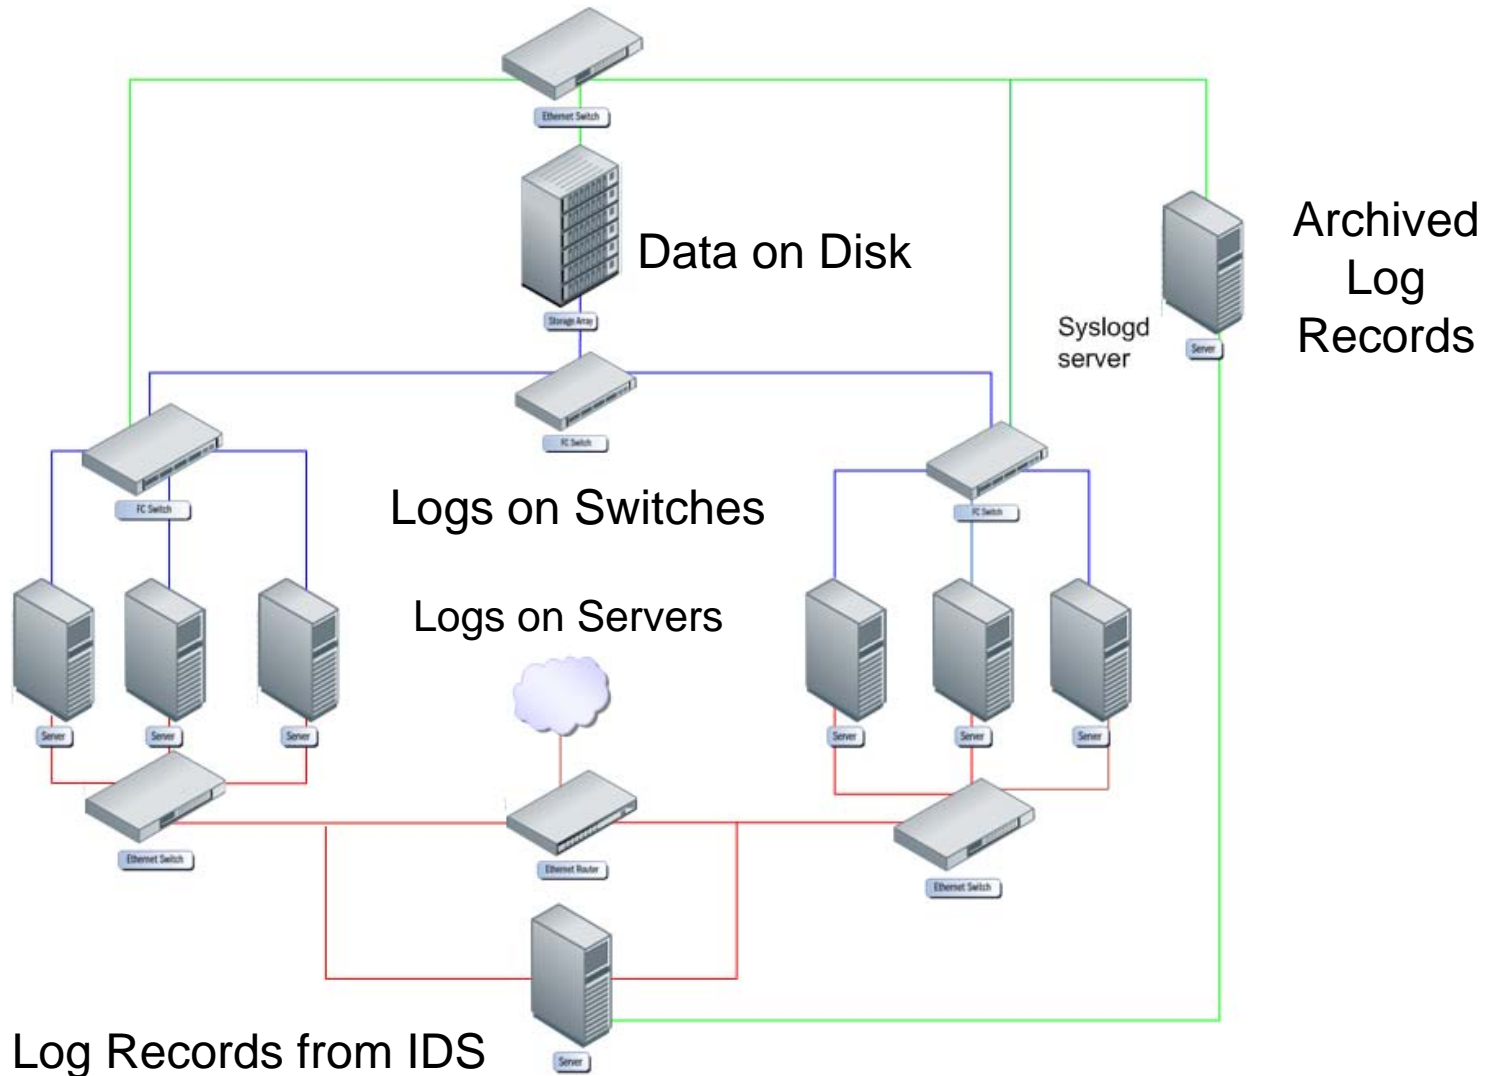
Banners

Management reserves the right to examine all data stored in or transmitted by these systems.

... without prior notice, management reserves the right to examine archived electronic mail, private file directories, hard disk drive files, and other information stored on company information systems.

...workers must have no expectation of privacy associated with the information they store in or send through these systems.

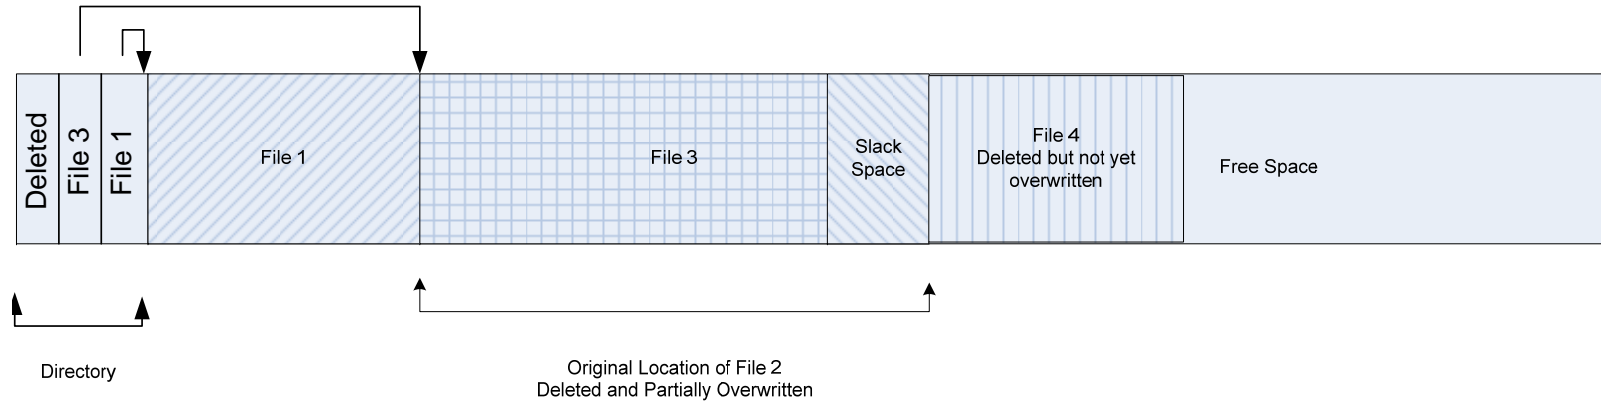
Sources of Evidence



Data on Disk

- Most common subject of forensic investigation
 - ◆ Also common in e-discovery for preserving information sources for later culling to identify relevant information
- Sound, repeatable process is critical
- Very common to create a bit stream image for later analysis

Bit Stream Image



A normal file copy would only copy File 1 and File 3

Un-deleting files before copying would retrieve file 4 but would modify the disk contents

A bit-stream copy that copied each block of the disk would retrieve the two files, the remnants of file 2 (slack space), the deleted file, the free space and the directory

- Digital information is highly volatile, easily changed or destroyed
- How can a specific piece of digital evidence be authenticated and its integrity verified?
 - ◆ By means of a “digital fingerprint” such as a cryptographic hash



```
@echo hello there>test.txt  
@md5sum test.txt  
d03d3fe9afff7a635879916173c1b383 *test.txt
```

```
@echo Hello There>test.txt  
@md5sum test.txt  
4e6a9dbf7699455525018b368c85d123 *test.txt
```

- Simply changing the case of two letters generated wildly different hash values
- Cryptographic hashes are used to produce a sort of “digital fingerprint” to demonstrate authenticity and integrity
- Commonly MD5 or SHA are used

Imaging Tools

- Wide variety to choose from:
 - ◆ Specialized Hardware tools
 - > \$\$\$
 - ◆ Software tools provided by forensic tool vendors
 - > \$\$
 - ◆ Standard Linux utilities such as dd
 - > Enhanced versions such as dcfldd and sdd
 - > FREE

- Become very familiar with the tool you use and make sure it is accepted in legal circles

Why A Linux CD?


- Linux is highly customizable and can be easily specialized for forensic use
- When dealing with a potentially compromised system, you cannot trust anything running on it
 - ◆ CD's are read-only
 - ◆ Are built with known versions of trusted tools

Imaging Walkthrough

- <http://www.e-fense.com/Helix> (free)
 - ◆ Downloadable ISO image
 - ◆ Both live response and bootable
 - ◆ Current version is 1.9a
- Introductory Guide
<http://www.e-fense.com/helix/Docs/Helix-for-Beginners.pdf>
- Optimized for forensic use
 - ◆ Specialized tools
 - ◆ Does not mount any disk filesystems
 - › Mounts a RAM-disk on /dev/ramdisk for temporary storage
- Other Options:
 - ◆ Knoppix STD
 - ◆ F.I.R.E.
 - ◆ Penguin Sleuth Kit



Field Imaging Setup



For SAN storage, a small server replaces the laptop

- Make sure the target disk is “sterile”
 - ◆ Put the soap away, that means it contains no traces of previous contents
- Calculate and record a baseline hash for the suspect disk
- Image the suspect disk
- Calculate and record the hash of the image
 - ◆ It should (and had better) match the one for the source
- Package image for transportation

dd command

- Used to generate a bit stream image
- Syntax
 - dd if=input_file of=output_file *options*
- Common options
 - bs=block_size
 - conv=noerror,notrunc, sync
- Example: Copy the contents of first IDE drive to the second IDE drive
 - dd if=/dev/hda of=/dev/hdb bs=4096 conv=noerror,notrunc, sync
- Note that target does not have to be the same size or technology -- target just has to be bigger than the source.
- `sdd` is an enhanced version with better performance

- Eliminate any traces of previous contents
- Prevents allegations of contamination
- Easily done by writing zeroes or random data to the target drive

Writing Zeroes

```
dd if=/dev/zero of=/dev/hdb bs=1024k
```

Writing Random Characters

```
dd if=/dev/urandom of=/dev/hdb bs=1024k
```

Imaging a Disk

The screenshot shows a terminal window with the following content:

```
Root Terminal
[root (knoppix)]# fdisk -l
Disk /dev/hda: 17.1 GB, 17179803648 bytes
255 heads, 63 sectors/track, 2088 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot      Start         End      Blocks   Id  System
   /dev/hda1  *          1         2087     16763796    7  HPFS/NTFS

Disk /dev/hdb: 41.9 GB, 41942138880 bytes
255 heads, 63 sectors/track, 5099 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot      Start         End      Blocks   Id  System
   /dev/hdb1          1         5099     40957686    83  Linux

[root (knoppix)]# mkdir /mnt/target
[root (knoppix)]# mount /dev/hdb1 /mnt/target
[root (knoppix)]#
```

Annotations on the screenshot:

- Suspect Drive**: Points to the HPFS/NTFS entry for /dev/hda1 in the fdisk output.
- Target Drive**: Points to the Linux entry for /dev/hdb1 in the fdisk output.
- Mount a filesystem on the TARGET so we can write files to it**: Points to the `mount /dev/hdb1 /mnt/target` command.
- Task is to image the SUSPECT drive /dev/hda to a file on the forensic TARGET /dev/hdb**: A large text block at the bottom of the terminal window.

The terminal window is overlaid on a blue background with a circular logo containing the text "ELECTRONIC DISCOVERY" and "PHYSICS". The system tray at the bottom shows icons for CPU, memory, network, and disks (hda, sda), along with the time 9:23 AM.

Imaging a Disk

- Calculate the baseline hash and store it as `b4HashItem01.MD5`
 - ◆ `md5sum /dev/hda>/mnt/target/b4HashItem01.MD5`
- Image the suspect disk to a file on the target as `Item01.IMG`
 - ◆ `dd if=/dev/hda of=/mnt/target/Item01.img bs=4096`
- Calculate the hash of the target and store it as `afHashItem01.MD5`
 - ◆ `md5sum /mnt/target/Item01.IMG>afHashItem01.MD5`
- Compare the hashes to verify they match
 - ◆

```
cat *.MD5
e70fb5d596d6544ad9a87e54f5928751 /dev/hda
e70fb5d596d6544ad9a87e54f5928751 /mnt/target/Item01.IMG
```

- Documentation helps assure both authenticity and integrity
- Create a permanent record of your actions
 - ◆ Permanently **BOUND** record book
 - ◆ Sufficient detail that a similarly experienced person could follow the same procedure and expect similar results
 - ◆ Sufficient detail to support your testimony even years after the actual event
 - ◆ Sign, date and have it witnessed
- Document in **WRITING** the hash values!
 - ◆ A signed, witnessed statement of these values supplements the electronic copies

Package and Secure

- Forms help you collect the relevant information
- Use tamper evident packaging
 - ◆ Seals “self destruct” when opened
- Pay attention to ESD and other precautions appropriate to magnetic media
- Store in a controlled-access location
 - ◆ A locked desk drawer will do if that’s all you have

- EVIDENCE -

(TO BE OPENED BY AUTHORIZED PERSONNEL ONLY)

Submitting Agency: _____

Case No.: _____ Item No.: _____

Date of Collection: _____ Time of Collection: _____

Collected By: _____ Badge No.: _____

Description of Enclosed Evidence: _____

Location Where Collected: _____

Type of Offense: _____


Victim's Full Name: _____

Suspect's Full Name: _____

Bag Sealed by: _____ Badge No.: _____

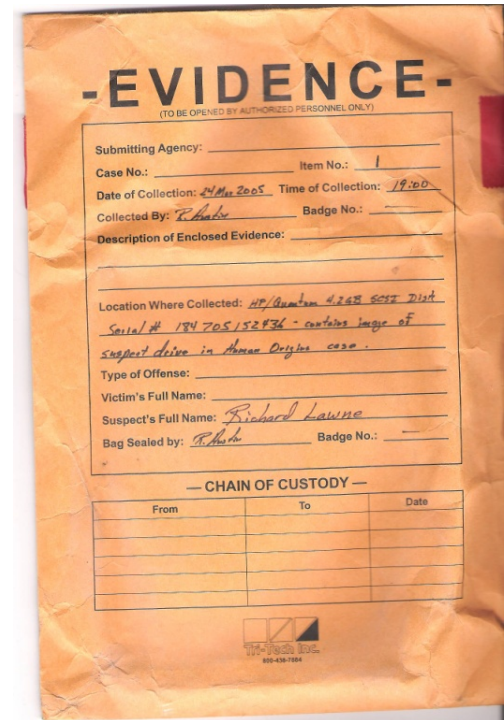
— CHAIN OF CUSTODY —

From	To	Date


 Tri-Tech Inc.
 800-438-7854



Package and Secure



Package and Secure

- Sign across all seals
 - ◆ Prevents someone else from opening the package and replacing the seal without detection
- Document the item both on a label attached to it and on the external envelope
- Maintain ***personal*** custody and control of the item until it is locked away

Package and Secure

CHAIN of CUSTODY LOG

Case Number

Item Number

PAGE
_____ CONTINUED FROM _____

ITEM DESCRIPTION

WARNING Receiver's signature warrants that evidence seal was intact with no visible sign of tampering except as noted under Notes at time of receipt

Relinquished By	Date / Time	Received By	Date / Time	Notes
PRINTED		PRINTED		
SIGNATURE		SIGNATURE		
PRINTED		PRINTED		
SIGNATURE		SIGNATURE		
PRINTED		PRINTED		
SIGNATURE		SIGNATURE		
PRINTED		PRINTED		
SIGNATURE		SIGNATURE		
PRINTED		PRINTED		
SIGNATURE		SIGNATURE		

The Chain of Custody form documents each and every access to the item.

Note the statement on the integrity of the seal as each person receives it.

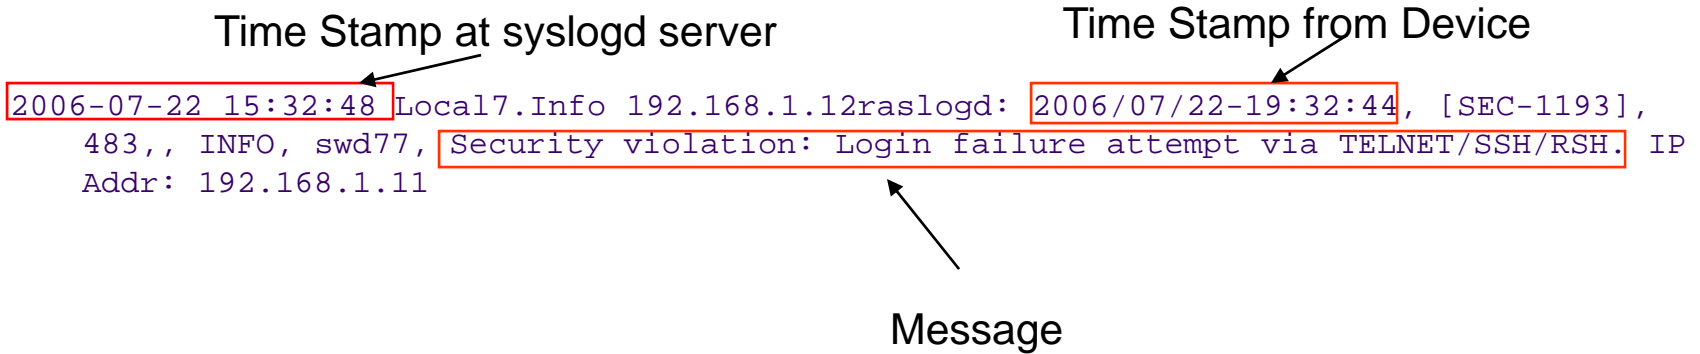
- Image the relevant LUN
 - ◆ A SAN LUN is really just a disk (as Roger Cummings once opined “just a bunch of sectors”)
- Use an imaging platform that supports fibre channel and allows you to mount the LUN as **read-only**
 - ◆ Linux-based tools work really well in this environment
 - › Make sure the distro supports your HBAs
- Once you have access to the LUN, image it just like it was a any other disk

- Forensic analysis typically handled by specially trained people or consultants using specialized tools
- These tools excel at dealing with deleted files, file fragments, operating system and browser artifacts

- Most devices common in the SAN environment are intelligent and generate event log records
 - ◆ A simple example is the authentication-failed messages a switch might generate during a brute force password attack
- These log records can be a valuable tool in establishing what happened and its progress over time

- Multiple sources of logs are a harsh reality in our environment
- Prepare in advance for being able to correlate log records
 - ◆ Accurate, standard time source such as a NTP server
 - ◆ Set switch timezone to “0” to use ZULU time
 - › While many switch vendors do support time zones, few automatically switch from “daylight savings” or “summer time” to standard time
 - › Avoid the whole problem with using accurate UT across the fabric

Example: Failed Logon



Note the 4hr difference between the server time stamp and the device time stamp (hint: this device is located in the US Eastern Time Zone)

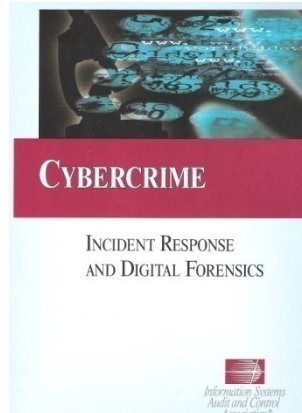






- Store logs off the device
 - ◆ Any “hacker” worthy of the name will attempt to destroy any record of their activities
 - ◆ For example, an intruder that modifies a zoning configuration to grant unauthorized access to LUNs will likely clear the switch log
- syslogd servers provide a centralized repository for log records in real time
 - ◆ See SNIA whitepaper “Audit Logging for Storage”
<http://www.snia.org/ssif/documents/SNIA-Logging-WP.050921.pdf>
 - ◆ See Guide to Computer Security Log Management (NIST SP800-92)
<http://csrc.nist.gov/publications/nistpubs/800-92/SP800-92.pdf>
 - ◆ Protect them appropriately (intruders know about syslogd too)

- The intruder may have cleared them (but this is an important fact as well)
- Will require live access to the device
 - ◆ Be prepared to testify to **EXACTLY** everything you did and its **consequences** to the information you collected
 - › For example, logging on to dump the logs will likely create a LOGON record
 - ◆ Do the **absolute minimum** necessary to collect the logs
 - › Resist any temptation to poke about to find out what the intruder may have done
- Download the log, immediately calculate a baseline hash and document it
- Treat the downloaded copy similarly to a disk image
 - ◆ Package and secure
 - ◆ Maintain a documented chain of custody

Centralized Logs

- “Freeze” (or snapshot) the logs
 - ◆ Establishes a definite point in time
 - ◆ Can be done by simply copying the log file (or more radically by pulling the network cable)
- Document what you do
 - ◆ Bound notebook
 - ◆ Witnessed record
- Use a cryptographic hash of the acquired log file to establish a baseline integrity reference
- Treat the copy of the log file just like it was a disk image
 - ◆ Package and secure
 - ◆ Maintain a documented chain of custody

For More Information



www.e-evidence.info

www.tritechusa.com – forensic supplies

<http://www.staticbags.com/>

- Please send any questions or comments on this presentation to SNIA:
tracksecurity@snia.org

**Many thanks to the following individuals
for their contributions to this tutorial.**

SNIA Education Committee

Richard Austin, CISSP

David Black

LeRoy Budnik, CISA

Roger Cummings

Eric Hibbard, CISA, CISSP

Larry Hofer, CISSP

Phil Huml

Michael Whitman, CISSP

SW Worth