

By FRANK ADELSTEIN

Live Forensics: DIAGNOSING YOUR SYSTEM WITHOUT KILLING IT FIRST

Live forensics gathers data from running systems, providing additional contextual information that is not available in a disk-only forensic analysis.

Fifteen years ago, servers for university computer science departments used disks with 2GB capacity; indeed, 1MB floppy disks were still useful. Many computers were connected to the Internet, many of them malicious, many at low, dial-up speeds of 9.6Kbps.

But, thanks to Moore's Law, that scale of operation is now considered meager. Today's desktop computers commonly have 2GB of memory, and more than 100GB of disk space. A 10-person company can have over a terabyte of disk space on its desktop machines alone. Home computer users connect to the Internet with high-speed links of 3–6Mbps, rivaling the bandwidth of smaller CS department networks not long ago.

TRADITIONAL DIGITAL FORENSICS ATTEMPTS TO PRESERVE ALL (DISK) EVIDENCE IN AN UNCHANGING STATE, WHILE LIVE DIGITAL FORENSIC TECHNIQUES SEEK TO TAKE A SNAPSHOT OF THE STATE OF THE COMPUTER SIMILAR TO A PHOTOGRAPH OF THE SCENE OF THE CRIME. FOR THESE REASONS AND MORE, INTEREST IN, AS WELL AS A NEED FOR, CONDUCTING A FORENSIC ANALYSIS OF LIVE SYSTEMS HAS ESCALATED.

The nature of digital forensic investigation has therefore changed. Larger disk capacities increase the time required for analysis and the difficulty and expense of collecting all disk evidence. The pervasive nature of the Internet makes contextual information more important—specifically, who is connecting to a machine and what they are doing.

In the traditional “snatch and grab” approach to computer forensics, an investigator pulls the plug on the machine, and then images (copies) the disk, either on site or (after confiscating the machine) in a lab. An analyst examines the image (actually a copy of the image) in a controlled environment—by repeatable steps on whose results both prosecution and defense will agree.

This approach has several drawbacks. First, it may not be possible. As systems continue to increase in size, terabytes of disk data are no longer uncommon and imaging can take many hours. Imaging is not possible, or at least extremely difficult on NAS, SANs, and large RAID arrays. The time and effort needed for analysis increases with the disk size. A person can quickly peruse all the sectors on a 1MB floppy disk. Manual scrutiny of a 1GB disk is not feasible; automated analysis assistance is required. Even automated analysis of a terabyte of data is slow [5].

A disk is offline while being imaged, but for many systems—such as e-commerce systems—the loss of

revenue from even a few hours of downtime is unacceptable. As a result, many judges no longer issue court orders to take down servers.

And, finally, much information about what is happening on a running system is lost when the plug is pulled. This information provides the *context* for the disk evidence. Traditional digital forensics attempts to preserve *all* (disk) evidence in an unchanging state, while live digital forensic techniques seek to take a snapshot of the state of the computer, similar to a photograph of the scene of the crime. For these reasons and more, interest in, as well as a need for, conducting a forensic analysis of live systems has escalated.

A warning and disclaimer: It is very easy to contaminate the evidence on a system. The many subtle points and quirks of various operating systems, as well as all the rules, constraints, and limitations imposed by the applicable legal system are beyond the scope of this article. Therefore any forensic investigation should be conducted by a professional.

WHAT INFORMATION IS AVAILABLE?

Information available from a live system provides a context for the disk data—for example, running processes, network connections, memory (process and physical), and other state items such as caches, logged-on users, and system load. Live analysis can capture both this volatile information and static information about the file system. Currently, most forensic tools use the operating system itself to obtain this information. If a machine has been com-

promised, its kernel can be malicious running code (a rootkit) that prevents the operating system from reporting the existence of processes and files. While there are programs that can detect the presence of rootkits (such as chkrootkit from www.chkrootkit.org), little can be done if a rootkit is present except to fall back to the traditional approach.

A big concern in live digital forensics is that the system is not static—files and processes are continually changing. However, this does not necessarily invalidate them as evidence. For example, even though the system log files will continually change and new mail continually arrives, this activity will not *create* an incriminating email message sent by the suspect days before the system is imaged. In addition, disk buffers may not have been written to disk due to caching. In other words, the context of the evidence is significant. The time stamps associated with the files, called “MAC time” (modify, access, and “create” times),¹ can help establish this time context.

ACQUIRING THE EVIDENCE

An investigation is an iterative process: repeatedly acquiring and analyzing data until a decision can be reached [1]. The investigator acquires the volatile state of the running machine by running programs on it. Since gathering evidence on the target can affect other evidence on the target, a set of best practices has evolved to maximize the quality of the evidence. Here, we discuss three of the most important practices.

Running known good binaries: An investigator should not trust the executables on the running system, but should provide all of the executables used for gathering evidence. The executables should be statically compiled, if possible. Otherwise, they should include any shared libraries required by the executable. The programs should originate from a read-only medium, such as a CD-ROM. The executables can be copied to the running system; however, this action will affect the disk, possibly overwriting evidence residing in deleted files. If the choice is between losing some evidence by overwriting and losing all evidence by not obtaining the information, it is better to risk the minor damage from copying files to the target system.

Hashing all evidence: Once acquired, evidence must be preserved in a way such that the investigator can later demonstrate that nothing has changed. The accepted method is to compute a cryptographically

secure hash of the data (typically via MD5 or SHA-1). The hash represents a fingerprint of the data with a small number of bytes, typically 16–20. The hash can be recalculated later and compared with the original to show the data has not changed from the time the original hash was obtained. If data is transmitted over a network from the target to another machine, the hash should be computed on both machines and compared to ensure no data changed in transit. Both the hashes and the evidence should be maintained in a secure location. Typically, an investigator preserves the integrity of the hash itself by signing and dating a printout of the hash and storing this in a secure location.

Gathering data in order of volatility: Some data is more ephemeral than others. Evidence should be gathered based on the Order of Volatility [4]. For example, open network connections change more frequently than the system load average or the users logged on to the system. Some actions may affect other data. For example, logging in to a system may generate entries in the system log files.

Complicating matters is the time required to gather evidence may depend on the kind of evidence gathered. A dump of the physical memory of a machine may be useful and is very volatile, indicating it should be accomplished early in the investigation. However, it can take tens of minutes to complete, and during that time, more useful information such as the lists of running processes, open files, and network connections, will have changed or disappeared. And while overall the system RAM is continuously changing, many memory pages may linger for a considerable time (days or weeks) on a modern system with 1GB or more of memory. In other words, the investigator must be aware of the overall context of the investigation in order to make informed decisions on the order of evidence acquisition.

PUTTING IT ALL TOGETHER

Once evidence has been acquired, it must be analyzed. In traditional forensics, the analyst gathers all potential evidence prior to analysis. It is often impractical to gather all possible information available to a live forensic examination. An investigator may therefore perform a triage and gather the essential data, examine it, and use the results of the first look to decide what else is needed. So for live forensics, the analysis step may lead to further acquisition of data. A live analysis creates opportunities for faster response.

Consider an example: The investigator receives a report of a slow Web server. He obtains a list of running processes and the network ports they have open.

¹The C time is sometimes referred to as the (inode) “change” time on Unix operating systems.

BECAUSE OF THE INCREASE IN THE QUANTITY OF DIGITAL EVIDENCE AVAILABLE IN EVERYDAY LIFE, IT WILL SOON BECOME IMPOSSIBLE TO ACQUIRE ALL DISK DATA RELATING TO A CASE. THE PARADIGM OF LIVE FORENSICS WILL BECOME THE ACCEPTED NORM.

One process has a connection to a high port number on an unknown system. The process also has a file open for writing. The file is a log of all network traffic, which indicates the process creating it is a network sniffer, most likely looking for personal information, passwords, among others.

This information could not be obtained by an after-the-fact analysis. In particular, the IP address and port of the unknown system could not be determined without a live forensic analysis.

CONCLUSION

Forensic data gathered from a live system can provide evidence that is not available in a static disk image. Live forensics also operates with different constraints—specifically, the evidence gathered represents a snapshot of a dynamic system that *cannot* be reproduced at a later date. Standards for acceptance are evolving, and legal precedents are still being established. An investigator faces risks that include evidence contamination and facing a court that does not understand the implications of the evidence. Because of the increase in the quantity of digital evidence available in everyday life, it will soon become impossible to acquire all disk data relating to a case. The paradigm of live forensics will become the accepted norm.

Live forensic data has been accepted in court cases, and tools exist to gather and analyze this evidence. But the field is still relatively new. Progress in several areas will be essential to increase the usefulness of live forensics, including tools to automate and standardize the process of evidence acquisition and preservation, and presentation tools that allow an investigator to present the facts clearly to a court.

Live forensics requires more effective memory analysis tools. Currently, it relies on the (possibly compromised) operating system to provide the list of

running processes. Live forensics needs tools to examine the raw memory of a machine and impose a process (and virtual memory) structure on the blocks of memory. These tools are analogous to the static tools that open the raw disk device and impose the file system structure on it to extract files, directories, and metadata.

Most memory analysis tools do little more than extract strings (ASCII for the naïve approach, Unicode for the more sophisticated). None imposes any process structure that provides application-specific information. Unfortunately, limited interest exists in decoding memory images because no one has done it, and no one has done it because the interest is limited. Gathering useful time-sensitive data and coupling it to fast, real-time analysis will provide new capabilities and insight for digital investigations and incident response. **■**

REFERENCES

1. Adelstein, F. The mobile forensic platform. In *Proceedings of the 2002 Digital Forensic Research Workshop*. (Syracuse, NY, Aug. 2002).
2. Carrier, B. *File System Forensic Analysis*. Addison-Wesley, Reading, PA., Mar. 2005.
3. Casey, E. *Digital Evidence and Computer Crime*. Academic Press, Mar. 2004.
4. Farmer, D. and Venema, W. *Forensic Discovery*. Addison-Wesley Professional, Dec. 2004.
5. Roussev, V. and Richard III, G.G. Breaking the performance wall: The case for distributed digital forensics. In *Proceedings of the 2004 Digital Forensics Research Workshop* (Baltimore, MD, Aug. 2004).

FRANK ADELSTEIN (fadelstein@atc-nycorp.com) is the technical director of computer security at ATC-NY in Ithaca, NY. He is the designer of the OnLine Digital Forensic Suite™ tool.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

© 2006 ACM 0001-0782/06/0200 \$5.00