IA 311: Network Forensics

Lecture 6: Investigating Network Equipment

Types of Wireless Access Points

- There are a wide variety of wireless access points available, General classes of WAPs include enterprise and consumer devices.
  - Enterprise

  Enterprise facilities typically span a much wider geographic range than home offices or small businesses.

  - Consumer

  Small businesses and home users often deploy consumer-class WAPs in their home and office environments. These devices are inexpensive and easy to configure for simple use

Why Investigate Wireless Access Points

Wireless access points are typically involved in forensic investigations for one of a few reasons:

- Wireless access points may contain locally stored logs of connection attempts authentication successes and failures and other local WAP activity.
- WAP logs can help you track the physical movements of a wireless client throughout a building or campus.
- The WAP configuration may provide insight regarding how an attacker gained access to the network.
- The WAP configuration may have been modified by an unauthorized party as part of an attack.
- The WAP itself may be compromise

WAP Evidence

- Wireless access points contain both volatile and nonvolatile evidence, although due to their persistent storage capabilities tend to be very limited.
- WAPs can also send logs over the network to a remote repository.
- As with switches and routers, most of the evidence on WAPs tends to be quite volatile.
- Enterprise-class WAPs tend to include the same functionality and range of evidence as wired routers, with the addition of wireless-specific capabilities
- Type of Evidence that can be fined on wireless access points:
  - History of connections by MAC address
  - List of IPs associated with MACs
  - Historical logs of wireless events access requests, key rotation, etc

Wireless Traffic Capture and Analysis

- Spectrum Analysis
- Wireless Passive Evidence Acquisition
- Analyzing 802.11 Efficiently

Spectrum Analysis

- There are, literally, an infinite number of frequencies over which data can be transmitted through the air.
- Sometimes the most challenging part of an investigator's job is simply identifying the wireless traffic in the first place.
- For Wi-Fi traffic, the IEEE utilizes three frequency ranges:
  - 2.4 GHz (802.11b/g/n)
  - 3.6 GHz (802.11y)
  - 5 GHz (802.11a/h/j/n)
- Each of these frequency ranges is divided into distinct channels, which are smaller frequency Bands.
- Although the IEEE has set globally recognized frequency boundaries for 802.11 protocols, individual countries typically allow only a subset of these frequency ranges

Wireless Passive Evidence Acquisition

- In order to capture wireless traffic, investigators need an 802.11 wireless card capable of running in Monitor mode.
- Many wireless cards do not support this capability.
- Furthermore, in order to ensure totally passive monitoring, it is preferable to use a special-purpose Wi-Fi monitoring card that can be configured to operate completely passively

Analyzing 802.11 Efficiently

- In order to analyze efficiently we can use tcpdump and tshark
- We can use Wireshark to sort out the endianness problem and for large packet captures in particular, tcpdump and tshark tend to be more efficient and scalable.

Why Investigate NIDS/NIPS?

- NIDS/NIPS alerts/logs may include details regarding illicit connections or even attempts. Can provide a richer source of information.
- NIDS/NIPS can be configured to alert on, or at least log traffic that firewalls deem perfectly acceptable.
- An investigator could potentially modify a NIDS/NIPS configuration to begin detecting events it wasn't previously configured to record.
- Rarely, the NIDS/NIPS itself might be suspected of compromisation

Typical NIDS/NIPS Functionality

- NIDS/NIPS typically features include:
  - Rules: Descriptions of how to compare a packet or stream with known malicious traffic.
  - Alerts: Lists of suspicious packets/streams.
  - Packet captures: Certain NIDS/NIPS can be configured to capture suspicious packets and save them for later analysis

Sniffing

- Passively: Connecting NIDS/NIPS to a mirroring port on a switch. Traffic is not detained.
- Inline: Putting the NIDS/NIPS itself inline between two devices in a choke point position. Can cause noticeable latency

Higher-Layer Protocol Awareness

- Protocol Reassembly
  - Attackers intentionally fragment the attack traffic. NIDS to perform fragment reassembly in order to be able to see and evaluate the target would receive and process.
- Normalization
  - HTTP protocol allows for many different ways of encoding data.
  - write a “signature” to detect each and every one of the various ways such a string can be represented.
  - normalize each string to its “canonical” form, and only then compare them against the list of “known bad” things

Alerting on Suspicious Bits

NIDS/NIPS can be configured to communicate alerts via various means. The most common of these include:

- Sending email alerts
- Logging events to a syslog server
- Sending SNMP traps
- Logging events directly to a queriable database
- Store alert and event data locally

Modes of Detection

- Signature-Based Analysis:

  This method compares headers, contents of packets, and streams of packets against databases of known, malicious byte sequences in order to identify suspicious traffic.

- Protocol Awareness:

  Protocol-aware NIDS/NIPS reassemble fragments (Layer 3), reassemble streams (Layer 4), and even reconstruct entire protocols (Layer 7) because they have to understand how the endpoint of the communication will interpret the data.

- Behavioral Analysis:

  The idea behind this approach is to spend time building a “baseline” of normal network behavior against which to compare future behavior. “learns” what is normal, in order to later detect aberrations

Types of NIDS/NIPSs

- Commercial
  - Check Point IPS-1
  - Cisco IPS4
  - Corero Network Security
  - Enterasys IPS6
- Roll-Your-Own
  - It is both free for use under the GNU Public License (GPL)
  - Supported by commercially funded research and development efforts.

NIDS/NIPS Evidence Acquisition

Types of Evidence

- Configuration
  - NIDS/NIPS has not been configured to alert upon a particular event, then the absence of alerts for event is meaningless.
- Alert data
  - Evaluate network traffic based on a set of rules or learned patterns, identify traffic of interest, and produce alerts.
- Packet header and/or flow record information
  - Log packet header and/or flow record data. help identify the origin, destination, and patterns of activity
- Packet payloads
  - NIDS/NIPS capture full packet contents from packets that trigger alerts, and subsequent packets.
- Activities correlated across multiple sensors
  - Topologically and temporally correlated event data is exceptionally useful

NIDS/NIPS Interfaces

NIDS/NIPS are almost uniformly designed to be information aggregation devices, so they are typically deployed with a central analysis console.

- GUI Interfaces
  - accessed, inspected, and configured via some sort of graphical user interface.
  - web client, special client
- CLI Interfaces
  - via SSH or direct console connection and use the command-line interface to view configuration, logs, and alerts.
- Off-System Logging
  - NIDS/NIPS systems are capable of sending critical evidence to an aggregation repository

Comprehensive Packet Logging

- Full content packet logging is occurring on NIDS/NIPS.
- Only packets being captured and made available for future inspection are trigger alerts.
- NIDS/NIPS generally are not configured to capture the packets that preceded an incident, nor the ones that followed.
- Problem for the investigator: often very little can be gleaned from a single packet, as too much of the context is unavailable

Storage Media

- The types of storage used on switches, routers, and firewalls vary between manufacturers and models.
- It is important for forensic investigators to be familiar with common types of storage used in network equipment in order to properly prioritize evidence collection.
- Understanding the volatility of data on different storage mediums is paramount, and as a general rule evidence should be collected and preserved in order of volatility (beginning with the most volatile first).

Storage Media …

Common types of storage in switches, routers, and firewalls include (in approximate order of volatility):

- Dynamic Random-Access Memory (DRAM)
  - DRAM is very volatile and does not retain data (for long) when power is turned off.
  - It is commonly used to store running operating system configuration, process memory, routing tables, firewall statistics, and more.
- Content-Addressable Memory (CAM)
  - CAM is a special kind of very fast memory used to store information that must be accessed extremely quickly.
  - It is most famously used on switches for storing tables that map MAC addresses to ports (hence the name “CAM tables”).
  - CAM is very volatile and does not retain data when power is turned off.

Storage Media …

- Nonvolatile Random-Access Memory (NVRAM)
  - NVRAM retains data when the power is turned off, but can also be easily modified.
  - The most common type of NVRAM found in network equipment is “flash memory.”
  - In routers, this often contains a copy of the operating system used at boot, as well as startup configuration files.
- Hard drive
  - Most switches, routers, and firewalls do not include a hard drive.
  - However, general-purpose servers can be configured to act as routers or firewalls (i.e., a Linux system running iptables).
  - In these cases, the hard drive typically contains the operating system, startup configuration, firewall logs, and an extensive amount of other data.
  - The data on a hard drive remains after the power is turned off.

Storage Media …

- Read-Only Memory (ROM)
  - ROM is a type of random-access memory that is designed to permanently store data without modification (hence the name).
  - ROM is not designed to be routinely modified, although nowadays types of memory commonly referred to as “ROM” can be reprogrammed in order to update firmware.
  - For example, on unmanaged switches, the operating system is typically stored in ROM.
  - For more capable and flexible managed switches and routers, the ROM typically contains a boot loader, which loads the operating system and configuration from NVRAM.
  - On fully configurable Linux systems that are used as routers or firewalls, the ROM normally contains the boot loader.

Why Investigate Switches?

Switches are typically involved in investigations for one of a few reasons:

- If you are trying to sniff traffic on a local segment, one of the easiest ways is to set up port mirroring on the switch.
- Switches contain tables that map client network card addresses (MAC addresses) to physical ports on a switch. This can help you to physically track down a computer.
- Attackers may launch attacks designed to “confuse” the switch in order to bypass network security restrictions or launch man-in-the-middle attacks. Forensic analysis of the switch may help to identify and isolate attacks of this type.

Switch Evidence

Volatile

- The majority of evidence on a switch tends to be volatile, stored in very fast memory such as CAM, or RAM for less time-sensitive applications. Volatile evidence on a switch can include:
  - Stored packets before they are forwarded
  - CAM tables—MAC address to port mappings
  - ARP table—MAC address to IP address mappings
  - Access control lists
  - I/O memory
  - Running configuration
  - Processor memory
  - Flow data and related statistics

Persistent

Switches also include persistent storage media that contain evidence such as:

- Operating system image
- Boot loader
- Startup configuration files

Off-System

- Managed switches (and some smart switches) may support automatic exportation of log data or flow data.
- Review the configuration files and architecture documents to identify these remote storage systems and retrieve the data.
- In addition, backup copies of configuration files are often stored on an external TFTP server.

Why Investigate Routers?

Routers are typically involved in investigations for one of a few reasons:

- Traffic of interest may traverse the router, resulting in associated flow data and related records.
- The network topology is key to understanding evidence and incidents, and is described at Layer 3 by the aggregate of routing tables.
- The router itself may be compromised.

Router Evidence

Volatile

As with switches, much of the evidence on a router tends to be volatile. This includes:

- Routing tables
- Stored packets before they are forwarded
- Packet counts and statistics
- ARP table—MAC address to IP address mappings
- DHCP lease assignments
- Access control lists
- I/O memory
- Running configuration
- Processor memory
- Flow data and related statistics

Persistent

Commonly, the following files are maintained in persistent storage:

- Operating system image
- Boot loader
- Startup configuration files
- Access logs, DHCP logs, etc. (primarily for roll-your-own routers with hard drives)

Off-System

- Routers tend to include very little, if any, writable persistent storage on board.
- However, most enterprise-class devices can be configured to automatically export data to external systems for storage, through syslog, FTP, TFTP, SNMP, and other means.
- Review the capabilities of the specific make and model you are investigating, as well as the running configuration information, to determine what, if any, data is being exported for storage elsewhere.
- Review of external storage media can provide you with a wealth of information such as the router access history, DHCP logs, backup configuration, flow records, and more.

Why Investigate Firewalls?

Firewalls often play a central role in network forensic investigations, for one of a few reasons:

- Firewall logs tend to include extensive information about connection attempts, whether or not these were successful, and if so, how much data was transferred from source to destination. Firewall logs may also include extensive details regarding protocols and applications in use, or even packet contents.
- Firewall configuration can reveal whether services or data were exposed to the world, or to systems of interest. It can also inform an investigator as to the type of evidence that logs may or may not include.
- An investigator may need to modify firewall configuration in order to configure the firewall to collect more evidence, or to gain access to systems of interest during the course of an investigation.
- The firewall itself may be compromised.

Firewall Evidence

- As with switches and routers, a great deal of evidence generated by firewalls is stored in volatile memory.
- “Roll-your-own” firewalls, which are built on top of a general-purpose operating system and hardware, tend to have more persistent storage available, since they are usually built with Commercial Off-The-Shelf (COTS) hard drives.

Volatile

Volatile evidence gathered from a firewall can include:

- Current, running configuration, including:
  - Interface configurations
  - ACLs and other firewall rules
  - Tunnels and their current states
  - Routing table and ARP cache
- Packet- and frame-level statistics
- Processor and I/O memory
- Command history

Persistent

Firewalls do store evidence in less volatile forms, including:

- Initial boot load
- Startup configuration
- Other static binaries
- Access logs, DHCP logs, etc. (primarily for roll-your-own firewalls with hard drives)

Off-System

- The evidence that any given firewall can produce varies by make and model.
- As a general rule, however, devices that include firewall capabilities tend to produce much richer logs and other data than more rudimentary appliances such as routers and switches.
- Evidence of an exported off-system may include, among other things:
  - Firewall logs
  - Access history
  - Backups of configurations

Why Investigate Web Proxies?

Web proxies are typically involved in investigations for one of several reasons:

- A user on the internal network is suspected of violating web browsing policy.
- An internal system has been compromised or may have downloaded malicious content via the web.
- There is concern that proprietary data may have been leaked through web-based exportation.
- A web server protected by a reverse web proxy is under attack or has been hacked.
- The web proxy itself has been hacked (rare).

Web Proxy Evidence

Persistent

- History of all HTTP or HTTPS traffic, including blogs, IM, web mail, etc.
- Blocked web traffic attempts
- Summarized user activity reports
- Web proxy configuration files

Volatile

- Cached content of web traffic stored in RAM.
- Cached content of web traffic stored on disk.
- Authentication information for web sites.

Off-System

- Web proxies may be configured to send system access logs, web surfing histories, and other records to a central log server.
- In some cases, large organizations may have a fleet of web proxies managed via a central console, with centralized logging and reporting.

Obtaining Evidence

- Web proxies often run on top of general-purpose operating systems, or versions of general-purpose operating systems that have been customized by a vendor.
- As a result, in many cases, the access options are similar to any general-purpose operating system, with the common addition of a web interface and/or proprietary interface as well.
- As forensic investigators, you may have the ability to collect:
  - Log files stored on the web proxy server or a logging server
  - Web cache files stored on the web proxy server
  - Reports from tools built into the web proxy server