ch #7
Issue
Detection
[part one]

LFS 260

Objectives

- Learn about ongoing issue detection across all attack surfaces.
- Deploy Suricata to detect network threats.
- Deploy OSSEC for server intrusion detection.
- Discuss behavioral analytics to counteract dynamic intrusion attacks.

Attack Life cycle

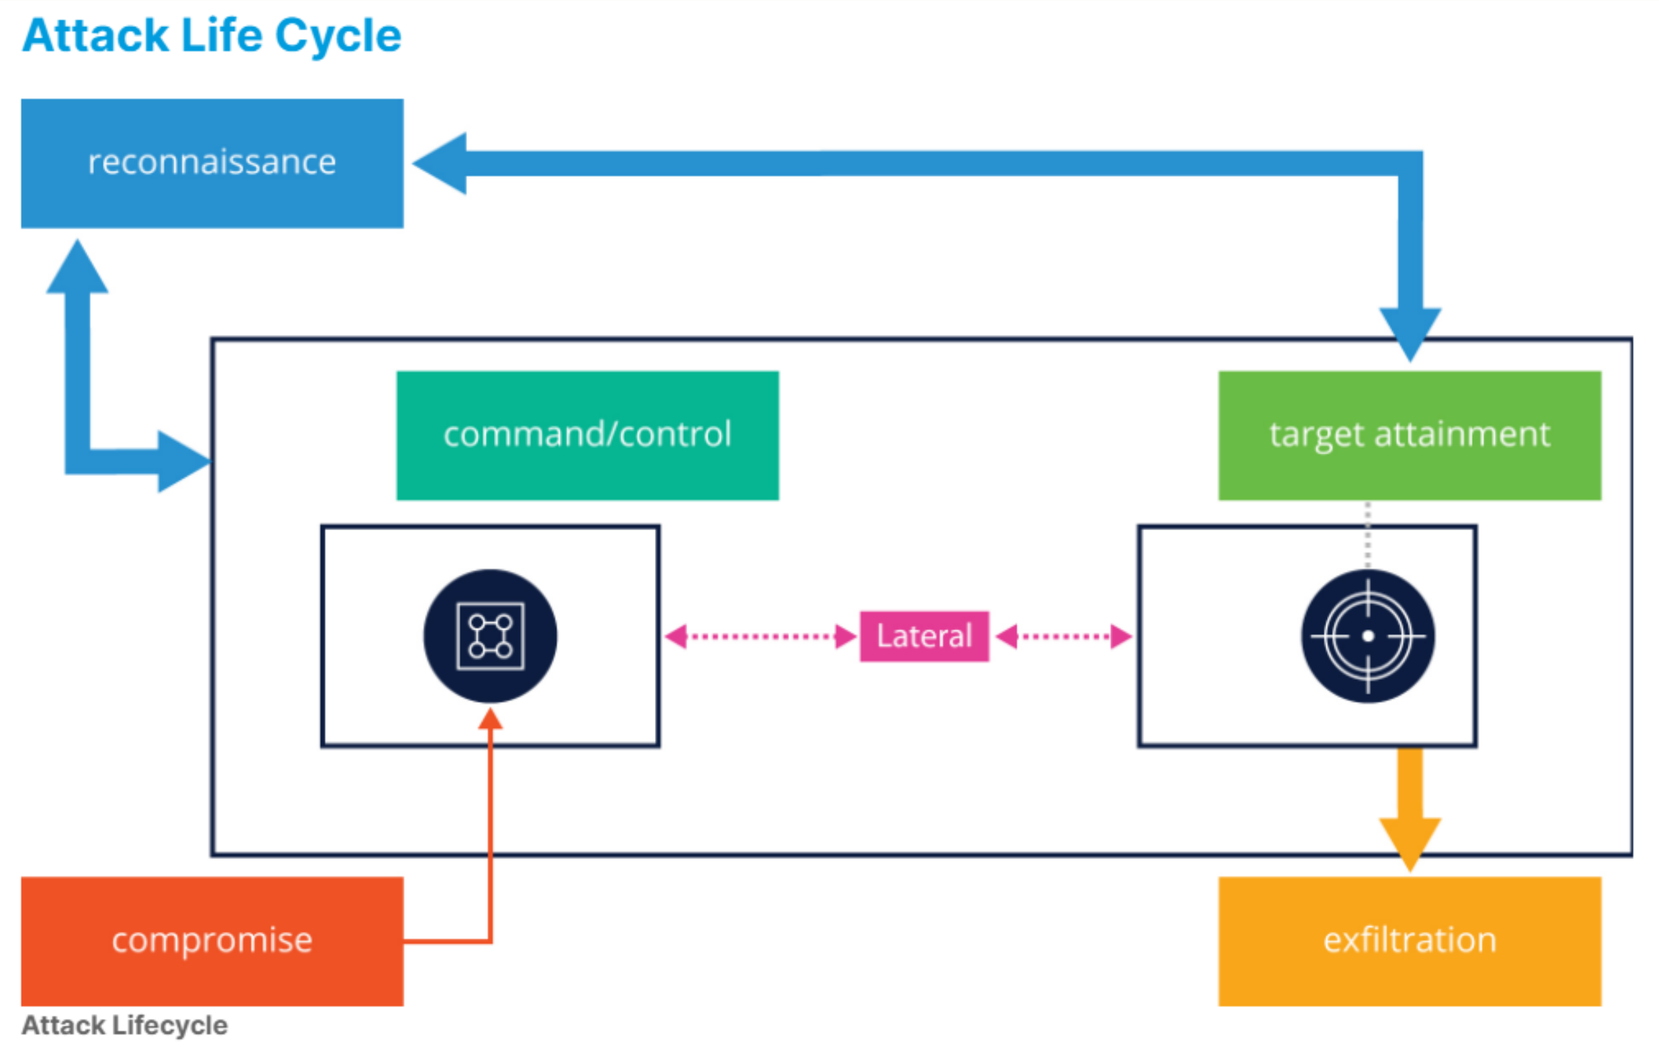

Attack Lifecycle

Attack Life cycle

During the **reconnaissance** phase the attacker will try to understand your environment as much as possible. The more information they can collect, the more they can determine a way to exploit your environment, looking for a vulnerability. This phase may range from a simple port-scan to buying drinks for IT staff and starting conversations about difficulties at work; anything to get more information about the environment. This phase could last years and can be difficult to detect.

With an idea of what exists, from hardware, to software, to personnel, the initial compromise phase begins by exploiting some weakness. This could be a **zero-day exploit**, or an ignored long-term vulnerability. This phase could last months, but if the reconnaissance phase went well, the attack should be focused and guaranteed to be quickly successful. If the attack does not work, the attacker will return to the reconnaissance phase. This phase is often the first time it’s possible to detect the attack.

Once a system has been compromised, the next phase is to take over control, and to create ways to continue control even if the intrusion is detected and the original compromise vector is closed.

Attack Life cycle

The **lateral movement** phase can be difficult to detect, as the attack takes over more and more systems. Done well, these steps may look like common administrative tasks, such as updating software or modifying configurations.

Once the valuable data has been located, the infrastructure has been documented, and multiple forms of compromise have been inserted, the target attainment phase completes. Remember that some compromises may be time-delayed, so that initial scans miss the eventual opening.

The final phase of exfiltration, corruption or disruption begins when the attacker disrupts or steals data.

Well Defined Security Policy

A comprehensive security/data breach policy is critical for any organization to ensure proper identification, mitigation, and response to a data breach, intrusion, or other attack and/or illegal action, both from external sources and from within (insider threat). Employee negligence and the insider threat should be weighed heavily when drafting any policy. This threat includes not only rogue employees, but simple negligence as well. A well-defined policy will enable a proper and measured response to both inside and outside threats.

Roles and Responsibilities need to be clearly defined, and a separation of duties is paramount to ensuring corners are not cut in the prevention, response, and mitigation of a data loss/breach.

Checklists should be created to ensure proper procedures are followed in the correct order and in a timely fashion. During an incident, confusion, panic, and disorder can be greatly reduced through proper preparation and a clear understanding of roles and responsibilities within an organization.

Keep a hard copy of the records in a secure location, both at the primary, as well as the recovery site. Technology companies often have mandates to keep all materials in electronic version, but if the system keeping the materials is also compromised, the recovery information could be changed or deleted.

Frequent Backups

Backups not only allow for disaster recovery/continuity of operations (DR/COOP), but greatly aid in the forensic investigation process. Backups both provide a snapshot in time, and can, in most cases, be submitted as legal evidence in a court of law. Backups should be stored both locally and off-site (consider 50 miles as a general rule) in case of catastrophic disaster.

Regular DR/COOP drills are essential to ensure both the readiness of personnel to respond to an incident, as well as discover any flaws or errors in both procedure, and hardware/software/media.

Even the most robust of RAID systems can be subject to failure. Any critical data on such systems should be part of your backup/DR plan.

Have regular restore drills. Too many organizations test once, then never revisit as their applications, hardware and staff continue to change. Instead, they hope their process and backups will work, only to find something missing or unusable at the worst time, during a major outage. While full scale recovery drills require a lot of time and effort, they are better than an extended outage, or inability to ever resume business.

Law Enforcement

Knowing who to call and when is a key component in any good security policy, but unfortunately it is also one of the more difficult questions to answer. Knowing when to notify law enforcement (local police, FBI, DHS, etc.) depends greatly on the jurisdiction in which your company lies (federal, state, local, military, etc.), and the severity of the incident. A good security policy will provide guidelines on how to measure the severity of an incident, and the steps for escalation.

Good questions to ask are:

- Does the agency in question conduct business across state lines, or on other continents?
- Is there classified data involved?
- If multiple agencies are involved, is there a required order for notification?
- Who is responsible to revisit the list of contacts and policy on a regular basis?

Doing this will greatly reduce the risk of evidence being corrupted, and aid greatly in measuring your response and understanding. As with backups, test the process. Ensure you have updated information, on a regularly scheduled basis.

Kill Chain

Kill chain, or cyber kill chain, is the attack progression. This refers to the steps taken by a malicious operator in order to compromise a target. Interrupting or preventing each step is an important part of cybersecurity.

Kill Chain

Reconnaissance

Research, identification and selection of targets, often represented as crawling Internet websites, such as conference proceedings and mailing lists, for email addresses, social relationships, or information on specific technologies.

Weaponization

Coupling a remote access trojan with an exploit into a deliverable payload, typically by means of an automated tool (weaponizer). Increasingly, client application data files such as Adobe Portable Document Format (PDF) or Microsoft Office documents serve as the weaponized deliverable.

Delivery

Transmission of the weapon to the targeted environment. The three most prevalent delivery vectors for weaponized payloads by APT actors, as observed by the Lockheed Martin Computer Incident Response Team (LM-CIRT) for the years 2004-2010, are email attachments, websites, and USB removable media.

Kill Chain

Exploitation

After the weapon is delivered to the victim host, exploitation triggers the intruders' code. Most often, exploitation targets an application or operating system vulnerability, but it could also more simply exploit the users themselves or leverage an operating system feature that auto-executes code.

Installation

Installation of a remote access trojan or backdoor on the victim system allows the adversary to maintain persistence inside the environment.

Command and Control (C2)

Typically, compromised hosts must beacon outbound to an Internet controller server to establish a C2 channel. APT malware especially requires manual interaction rather than conducting activity automatically. Once the C2 channel is established, intruders have hands on the keyboard access inside the target environment.

Kill Chain

Actions on Objectives

Only now, after progressing through the first six phases, can intruders take actions to achieve their original objectives. Typically, this objective is data exfiltration which involves collecting, encrypting and extracting information from the victim environment; violations of data integrity or availability are potential objectives as well. Alternatively, the intruders may only desire access to the initial victim box for use as a hop point to compromise additional systems and move laterally inside the network.

Emergency Response Team

Response teams should be well organized with roles and responsibilities clearly defined. Technical team members should have jump bags with any reference material, hardware, policies/checklists (hard copy), etc., prepared and on hand, ready to go at a minute's notice. Non-technical members should have (at the very least) a binder (hard copy) prepared with checklists, policy details, contact information, etc. As much as this information is essential for continuing operation, it also must be protected, as it contains a lot of sensitive information.

Regular drills should be conducted with performance expectations clearly stated, and if possible, a peer-defined grading scale such that constructive and somewhat objective feedback can be provided by senior management and staff.

Post Incident Forensics

Immediately following a security incident or data breach, a snapshot of the system in question should be created on write only media, or at the very least flagged as read only following creation. Any disparity in file/filesystem metadata can taint evidence.

When performing forensic analysis, a live Linux CD/DVD is the preferred method for accessing data from the image in question. The image should be mounted read-only, as care needs to be taken to preserve the integrity of the image. If investigators require a writable image of the compromised system, make a copy of the image, making sure to carefully label all media with consideration towards the documented chain of custody.

If an external intrusion is discovered, and the system or network in question is not mission-critical, consider disconnecting the system from the network, or even the network gateway if a more wide-spread issue is discovered. This will aid investigators in locating the source of the attack/breach, as well as go a long way toward limiting the damage that can be done if from an outside threat.

Host Intrusion Detection System

Host Intrusion Detection Systems typically are installed or scan for anomalies on the local node. Attributes such as MD5 sums, time-stamps, permissions, etc., are checked for anomalies. This requires a known good state to be recorded and kept in more than once place to avoid an attacker updating the attributes after changing important files.

AIDE (Advanced Intrusion Detection Environment) is shipped with Enterprise Linux and can be installed on all Linux operating systems. It is a highly configurable HIDS that requires an initial baseline database be created and then stored off the system for future validation.

Tripwire is similar to AIDE, with the addition of a commercial management console and real-time auditing agent. They do have an open source branch that has not had a new release in years, which can be found on GitHub.

Host Intrusion Detection System

OSSEC is an open source host-based intrusion detection system that performs log analysis, file integrity checking, policy monitoring, root-kit detection, real-time alerting, and active response. The OSSEC+ is available if you are willing to register, and provides extra features. Atomicorp also sells a feature-filled Atomic Enterprise OSSEC version of the tool. This tool uses a Central Manager which collects information from lightweight Agents capable of running on a wide range of operating systems. Some information can also be collected from Agentless systems and syslog messages generated by firewalls, switches, and routers as well.

Network Intrusion Detection System

Network Intrusion Detection Systems are tools that generally collect traffic from networking devices, typically routers, switches, and firewalls, and then analyze that traffic for attack signatures and other anomalies.

SNORT was one of the first such open source applications to fill this role. Several tools have been developed around integration with snort or directly including snort's functionality. The Cisco company develops, maintains, and provides commercial support for snort. Cisco Talus develops, tests, and approves rules which become part of the Snort Subscriber Ruleset, and require a subscription. Typically, these rules become part of the Community Ruleset after thirty days. You can learn more on snort.org.

Suricata is a newer NIDS tool, developed by the non-profit Open Information Security Foundation. According to their website

Physical Intrusion Detection System

Physical Intrusion Detection Systems have more to do with systems that ensure only authorized users are able to access the systems and environment. Consult with a physical security specialist when building the data center and workstation environment that has access to the data center.

Physical barriers such as fences are the first step. The more distance between public traffic and sensitive systems the better. A system of cameras and motion detectors should also blanket the exterior to allow constant monitoring. Every door should have a lock. While cards or RFID could be used, other means such as retinal scans or a combination lock could also be used. Ingress and egress should also be controlled and limited, in case someone sets a fire and takes a stack of hard drives while everyone exits. Windows are also a point to address, as they can be used to see inside the environment, as well as a hard surface for listening. Some facilities will use piped-in music and small motors to vibrate any required windows. Fewer is better.