INTRODUCTION TO CYBER DECEPTION

FOREWORD

What is cyber deception? Is it indeed a new paradigm in information security, or is it just about honeypots making a comeback?

For those who have been hearing the term cyber deception for some time now, but are still confused as to what it entails, this paper will introduce the concept and explain what deception is all about. It will also describe some of the pain points and holes that currently exist in cyber defense, and how deception addresses those issues. Lastly, it will explain why and how cyber deception can – and most likely will – be an important part of network security strategies in the future.

OUTLINE

This whitepaper will attempt to give a quick but comprehensive overview of cyber deception by covering the following topics:

- What is and is not typically considered as cyber deception?
- A primer on deception technology
- The problems and use cases that are best addressed using cyber deception, along with a deeper dive into sample deception architectures
- How to create and maintain effective deception campaigns that will correspond with organizational security goals

Cymmetria

© 2016 Cymmetria Inc.
All rights reserved. Confidential and proprietary.

TABLE OF CONTENTS

INTRODUCTION

Today, perimeter breaches are considered imminent. Once an attacker gets past a certain point in perimeter defense, defenders have little to no visibility of what is happening and thus have difficulties identifying and stopping attackers. Meanwhile, attackers are moving laterally within the organization or lying in wait for a good time to act.

Cyber deception is about baiting, studying, investigating, fingerprinting, and/or smoking out these attackers.

Deception catches threat actors as they make their first movements inside a network, by leveraging the fact that after gaining access to a network, attackers follow a predictable attack pattern: reconnaissance, lateral movement, and exploitation. Starting from the initial reconnaissance phase, deception technology takes advantage of this and creates a controlled path for attackers to follow. Deception focuses on critical stages through which all attackers must pass: infiltration and lateral movement. This allows for the hunting and catching of APTs, as well as less sophisticated threat actors.

Security is often an economics game; an attacker must consider how much to invest and what exactly their potential gain is. Successful deception shifts the dynamics of cybersecurity. Attackers no longer enjoy a situation in which they can move freely within a network, and return to the same network multiple times – often with the same exploits and tools. Instead, they are now forced to invest more resources, time, and effort in their attack attempts, and are under the constant fear that they might make the wrong move. In other words, deception creates an unfriendly environment for attackers, one in which running tools or exploits on the wrong target means the end of the attack, as attackers are fingerprinted and signatures of their attacks are generated and distributed throughout the organization.

CYBER DECEPTION – HONEYPOTS WERE UNSUCCESSFUL … WHAT CHANGED?

One of the most common reactions from people who hear the term cyber deception is “Oh, so it's like a honeypot!” Honeypots are definitely a component of cyber deception, but deception is not just about creating a machine that will look believable (i.e., an emulation) to an attacker, and then having that machine learn everything possible about said attacker.

As you will see in the next section, cyber deception is comprised of multiple elements – honeypots being just one of them. One of the reasons honeypots aren’t in widespread use today is that honeypots are simply not effective enough for mass usage; they were and are still considered to be very complex to deploy and maintain, and/or easy for attackers to detect. So, what is the difference between deception and honeypots?

Other than the fact that deception is made up of much more than just honeypots, there is also the fact that technology has evolved; orchestration and virtualization have now been commoditized, allowing for many things that were previously impossible. For example, even five years ago, if you wanted to set up a network of honeypots, you needed to manually distribute the honeypots among different servers, or find another way to make them look believable. If you wanted to create a honeypot on demand, you had to build the scripts. If you wanted to build a set of honeytokens and then automatically deploy them across the organization's network, there was simply no way to do so.

With current virtualization and orchestration technologies, it is now possible to:

1

build deception-on-demand, which adapts to attacker activities

2

centrally manage deception elements

3

automatically build deception into new systems that are being created in the organization's networks

Additionally, the new generation of cyber deception now allows for the integration of deception elements into an organization's existing security tools, where they can harness the power of these tools. As a result, the tools themselves become more valuable because of the additional information received from deception elements.

CYBER DECEPTION ELEMENTS – CREATING ALTERNATE REALITIES FOR ATTACKERS

Some of the elements of deception, such as honeytokens, have been around for a while. With the advent of cyber deception, these elements can now be used in new ways in order to stop APTs and attackers in the lateral movement phase, while also addressing other immediate needs. Some of the most common elements of deception include:

HONEYTOKENS, WHICH CAN TAKE SEVERAL DIFFERENT FORMS:

Breadcrumbs - These are pieces of data that are meant to be picked up by attackers who are gathering information necessary for lateral movement. These are placed everywhere, or at the very least in the critical path of attackers.

EXAMPLE

Cached RDP connections, Windows credentials, VPN configuration files, SSH private keys, browser history and cookies, and mapped network shares. All of these are likely to be encountered and used by an attacker who has gained a foothold on an endpoint.

Elements that phone home when touched - These are meant for detection purposes. The concept behind these elements is that they are not part of the routine operation of the organization, so if anyone touches them, it is almost certainly malicious activity. Of course, the elements have to be built in such a way that they will not be triggered by standard network activity.

EXAMPLE

Dummy network shares that are set to raise an alarm if someone touches them, documents that contain macros that send commands, and websites that send information to a SIEM if someone tries to navigate to them.

Tokens that are placed in the organization and are used to identify leaks.

EXAMPLE

Honeydocs, which are documents with external assets or DRM that, as soon as they are opened, act as a sort of beacon to alert the system of exfiltration.

DECOYS/HONEYPOTS

These are the attack targets in the deception story. Either breadcrumbs lead the attackers to them, or these targets are encountered randomly or statistically by the attackers. These decoys are designed and built to resemble attacker targets.

Decoys perform a variety of functions:

- Slow attackers by providing a fake target, on which attackers waste time.
- Shift the burden of anomaly detection from defender to attacker, forcing the attacker to be on constant alert in order to detect decoys and other fake elements in an organization's network. If attackers know of or take into consideration the existence of decoys, they have to be much more careful in their attack attempt.
- Gather extensive information about what is being done to them. This allows decoys to serve as a sandbox, providing an analysis of everything happening inside. This includes samples of files that were downloaded or executed on the decoy, a memory dump during an exploit or file execution, a network traffic dump, the history of attacker-executed commands, the address of outbound connections from the attacked machine, and more.

DECEPTION STORIES

The deception story is the combination of breadcrumbs and decoys portrayed in a believable manner that matches the normal business processes of an organization. Deception stories are essentially the shaping of deception elements into scenarios.

HOW DECEPTION FITS WITH ORGANIZATIONAL SECURITY STRATEGY

The following scenarios are the most common instances in which cyber deception is used.

STOPPING AND DETERRING ATTACKERS

Need

Stopping attackers who have managed to breach the organizational perimeter, by making the organization a difficult target for advanced threat actors, and providing high-fidelity alerts when deception targets are touched. Another goal is to deter attackers and cause opportunistic attackers to seek other targets.

Solution

The most fundamental aspect of a good deception story is its ability to lead attackers down a controlled path, and divert them from real organizational assets. This does not mean attackers will necessarily overlook real assets, but strong deception will allow defenders to stop the attacker or closely monitor their actions even if they access actual organizational targets along with deception elements.

If attackers are aware that a deception system has been deployed throughout an organization's network, that knowledge instantly changes their behavior. This causes attackers to move more slowly and plan each step more carefully, since a mistake would land them inside the deception network. Executing their tools on the wrong system would result in the defender gaining information that would enable them to analyze the attacker's behavior, and understand their tools as well as their modus operandi.

RAPID DETECTION

Need

Drastically reducing attacker detection time requires immediate, high-fidelity alerts that do not need processing or analysis.

Solution

Deception elements are not a routine part of the organizational network and therefore are not susceptible to being set off by legitimate users during their normal daily activity. If an alert is triggered, it means with nearly 100 percent certainty that an attacker is at work.

It is possible to feed Windows credentials breadcrumbs into the organizational SIEM; if these are stolen by an attacker and used anywhere in the domain, the failed login attempt will trigger an alert.

Similar to breadcrumbs, fake SMB shares or websites can operate in the same fashion. Since they are not listed anywhere and are not meant to be used, if they are accessed it means that someone is doing something they should not be doing.

The most critical element here is accuracy and avoiding false positives and alert fatigue. The functionality of this use of cyber deception depends entirely on the fidelity of generated alerts. Deception elements such as decoys can deliver quite verbose alerts (for example, every time a port on a decoy is scanned), but they can also be tuned to deliver only the most relevant alerts, such as the execution of code on one of the decoys. This drastically reduces the number of alerts while still maintaining visibility of attacker activities.

TECHNICAL INFORMATION

On Windows machines, there is a utility called cmdkey. This utility allows the user to add stored credentials to the machines. By fabricating user and password combinations and adding them to the stored credentials on an endpoint, we create a mechanism for detecting when the system has been breached. The best way to detect the usage of a specific user/password combination is by adding a rule in the SIEM and creating authentication events that look for that particular user/password combination. As soon as anyone uses that combination to try to authenticate themself on another machine, it is obvious that the original endpoint has, by definition, been breached.

SUPPLY CHAIN PROTECTION

Need

It is often difficult to control the security of the supply chain, which consists of vendors, suppliers, and companies that an organization has acquired and needs to onboard quickly.

Solution

In order to address this problem, the organization can cast a deception web over supply chain elements as well. For example, during supplier onboarding, the supplier will access a web portal that places a cookie. Attackers following this cookie will reach a deception element that will trigger an alert.

TECHNICAL INFORMATION

One option for casting a net of protection over suppliers is requiring VPN use. During onboarding, an organization can provide suppliers with two VPN configuration files: one of these files is to be used when actually accessing the network, while the other is meant to be left aside for attackers to find.

Another option is placing domain credentials and using rapid-detection mechanisms (as detailed in the rapid detection section of this paper). Specific policies can be set for suppliers who have access to the domain and deception breadcrumbs can be placed on their machines.

An additional option is placing a browser cookie in the browsing history of suppliers who visit an onboarding portal. This cookie, when detected by attackers, points to a deception domain. This works for remote branches and new companies that have been onboarded.

INTERNAL THREAT

Need

Malicious insiders are hard to detect since they are privileged and are subject to fewer security controls. As a result, there is a need for some kind of effective identification of insider threats making steps inside the network.

Solution

It is necessary to generate a type of deception that is attractive to insiders. Some possible options include setting up SMB shares, placing honeytoken documents, and inserting fake records into the database.

TECHNICAL INFORMATION

When confronted with several open permissions for network-shared folders, an attacker will eventually look into those folders. These kinds of deception elements can be created using the "net use" command on Windows machines, and when not mapped to a drive letter, will not appear in the Windows GUI in any way. These elements can only be accessed by running the net use command or by directly giving their network address or UNC. This is important because regular users should not be able to trigger this by going into the mapped drive, and yet it makes sense that a specific shared folder would appear in the net use list, since folders are added to the list once they have been visited.

AUGMENTING THREAT INTELLIGENCE, SANDBOXES, INCIDENT RESPONSE, AND OTHER SECURITY TOOLS

Need

Many security tools are difficult to make effective, or require a lot of money and resources. They generate a lot of information, such as alerts and raw data.

Solution

Augmenting security tools with high-fidelity information allows harvesting more value out of them. The breadcrumb and decoy elements of deception allow organizations to effectively sandbox an attacker and the code or malware the attacker executes.

When attackers execute commands on a decoy, they are running inside a sandbox/honeypot. Attacks and malware executed against a decoy will lead to the fingerprinting of the attacker and their malware; forensic data is gathered and signatures are created. This allows immunization against the malware, as well as the identification of zero-days.

Another aspect of the incident response is the possibility of identifying the attack origin on a breached endpoint, by placing only unique breadcrumbs on endpoints. Additionally, by rotating breadcrumbs after a predefined period of time, it is possible to pinpoint the time of breach.

When considering threat intelligence and the SOC, it is important to remember that any information gathered by deception elements is high fidelity, and as such, can be used either by tools that rely heavily on event scoring or by tools that are built to gather and process as much information as possible.

DECEPTION MATURITY AND MAINTAINING EFFECTIVE DECEPTION STORIES

In order to facilitate the adoption of a deception solution, it is best to have one that is easy to use and has a quick initial rollout; however, its ability to mature with the organization is important and should also be considered.

Some deception solutions come with pre-defined deception stories, which consist of one-size-fits-all elements that are likely to be found in most organizations, such as intranet websites or database servers. After those have been deployed, it is possible to create more elaborate and convincing deception stories that will target the more advanced attackers.

PREPARATION

THREAT ASSESSMENT AND SECURITY STRATEGY

Deception deployment depends on the issues the organization wishes to solve. In order to determine the structure of the organization-specific deception, the best method is to look over the use cases described in the previous section, to see which of them aligns with the organizational security focus.

It should be possible to deploy pre-defined deception stories that address the threats that are most relevant to the organization, and then deploy more tailored stories that match specific organizational concerns and needs.

TECHNICAL PREPARATIONS

One important preparation is network mapping. This is necessary in order to identify in which VLANs the decoys will be deployed (often all of them), the location of the endpoints that need protection, and which types of breadcrumbs and decoys need to be deployed. Most modern cyber deception solutions are centrally managed, so a management interface should be placed in a location accessible from all decoys/honeypots. Existing security tools should also be mapped in order to be connected with information coming from the deception system when possible.

CRITICAL ASSETS

Critical assets that are likely to be targeted by advanced attackers should be mapped out as well. This will ease the process of placing convincing breadcrumbs in the path of attackers who are navigating their way toward these assets, and result in the creation of believable decoys that can resemble these assets and lure attackers into interacting with them.

DEPLOYMENT

The deployment stage involves creating a deception scenario based on the information gathered in previous stages. This is essentially the strategic creation of breadcrumbs and their subsequent placement at attacker entry points (in most cases, this means all the endpoints and servers in the network, as it is never possible to know where an attack will come from). Another critical part of the process is the creation of decoys that are relevant to the organizational architecture and structure, as well as the critical assets these decoys are meant to protect.

To ensure the deception system is properly functioning, periodical health checks are recommended, particularly in systems that are generating few alerts during standard operation. An example of a system check could be a weekly status email detailing deception system activity.

Also, the deception system launch should be followed directly by a penetration-testing project. Some of these projects and tests are performed with the penetration testing team having some kind of prior knowledge of the deception system, while others are not. This discrepancy helps test the effectiveness of both the actual deception and the penetration testers.

CONCLUSION

Deception technology is quickly gaining popularity due to the increasing need for an effective solution for stopping and deterring attackers, gathering quality attack analytics and information, and generating high-fidelity alerts.

While deception does involve honeypot technology, it is a complex field that meets multiple needs, many of which can only be addressed by deception. The effects of cyber deception are far reaching, and include benefits such as rapid attack detection, attack origin pinpointing, supply chain protection, the identification of advanced attackers, and more. Not only is cyber deception highly useful and powerful, but it can usually be introduced gradually to the organization without requiring many resources, which makes it a technology that is relatively easy to adopt. There is little doubt that cyber deception is a tool of the future and will soon become an integral part of cybersecurity systems the world over.

FOR MORE INFORMATION OR FOR A PRODUCT DEMONSTRATION, PLEASE CONTACT IRENE AT IRENE@CYMMETRIA.COM