Using Deception to Protect Microsoft Active Directory

Executive summary

This paper describes an example of how deception technology can be used to detect (and therefore control) an attacker as they attempt to breach your Active Directory installation.

- It covers the basic concepts behind a typical campaign deployed to protect an Active Directory installation using the CounterCraft Cyber Deception Platform and the infrastructure required to achieve the steps explained. The deception technology is used to detect this activity in three distinct areas; enumeration of AD credentials on endpoints; enumeration of AD credentials on the production AD domain controller; and enumeration of AD credentials in shared resources.

- The example concludes that the use of deception technology to protect an Active Directory installation requires the use of a carefully structured Deception Campaign, with some or all of the components described in the paper.

- In summary, Active Directory is such a critical component, that protecting it correctly is key. Using deception technology provides a way to add an extra layer of protection to your existing security systems, to detect, investigate and control any potential attack.

For more information, a comprehensive demo or to start the process for full implementation within your organization, get in touch with us directly at craft@countercraftsec.com.

Introduction

Microsoft Active directory is really the default enterprise network operating system. It's everywhere. It is where we store all our network, user and infrastructure data. To quote the O'Reilly big book of "Active Directory"[1]:

"*Active Directory enables administrators to manage enterprise-wide information efficiently from a central repository that can be globally distributed. Once information about users and groups, computers and printers, applications and services has been added to Active Directory, it can be made available for use throughout the entire enterprise…*"

Being a single, global repository for so much network information makes Active Directory (AD) a very obvious target. The purpose of this paper is to describe an example of how deception technology can be used to detect (and therefore control) an attacker as they attempt to breach your AD installation.

To begin with, let's set out the stall. Throughout this paper we will talk about **production** and **deception environment**s. To be clear, the **production environment** is your standard IT infrastructure. It is all the endpoints, servers and services that run in your organisation. The **deception environment** is a fake environment built by you with the sole purpose of enticing would-be attackers. The idea is to build a realistic and credible simulation, but not an exact replica of your **production environment**. It includes all the deception hosts and services created to spoof the attacker. In this example, we will make a basic assumption that an endpoint has been already compromised. This means functional domain credentials – a username and password – have been obtained by the attacker.

This second assumption is that the attacker will follow the most common attack-tree, or adversary activity map, for this scenario, which is as follows:

1) Compromise of a production endpoint – e.g. a user workstation.

2) The attacker then tries to identify users and credentials for the domain (this is also known as Local Domain Mapping).

3) The next step is to attempt to enumerate users, domains and policies from the domain controller (also known as LDAP reconnaissance).

4) Obtain credentials located in the local machine, also known as Credential Theft.

These are common steps taken by any attacker trying to compromise an AD installation. We are talking about the attack stages defined by the Mitre ATT&CK™ Matrix from "Execution" to "Collection".

Deception technology can be deployed to detect this activity in three areas:

> **1)** Detect Enumeration of AD Credentials on Endpoints
>
> **2)** Detect Enumeration of AD Credentials on the Production AD Domain Controller
>
> **3)** Detecting Enumeration of AD Credentials in Shared Resources

Every area is crafted as part of a carefully structured deception campaign. This is deployed and managed from the CounterCraft Deception Director.

The next section explains the different areas in greater detail.

1 Detecting Enumeration of Active Directory Credentials on Production Endpoints

To detect the enumeration of users and credentials, endpoints are seeded with specific breadcrumbs pointing to a fully instrumented deception AD Domain Controller (AD DC). Any attempt to interact with the deception AD DC will instantly be detected, and trigger notifications on the CounterCraft Deception Director, to alert on the malicious activity.

Additionally, deception user credentials linked to the production AD DC are also placed on local machines. When they are used, the false credentials create an alert from the AD. Assuming that the production AD DC logs are collected by a SIEM, the CounterCraft Deception Director will take a log feed from the SIEM to enable the generation of CounterCraft notifications. Alternatively, the AD DC can feed logs directly to the Deception Director.

The next part is to seed the endpoints with deception credentials to deception services hosted on additional deception servers, located in a deception environment. Any lateral movement to these resources will be instantly detected and alerts sent by the Deception Director.

2 Detecting Enumeration of Active Directory Credentials on a Production Active Directory Domain Controller

The production AD DC is seeded with fake users, credentials and resources. If the attacker decides to ignore, or has not found, the breadcrumbs that lead to the deception AD Domain Controller, and performs an enumeration of users, domains or policies of the production AD Domain Controller, these false credentials will be presented.

When the fake credentials are presented, an alert is raised from the AD DC. As before, this is typically reported via a log feed from the SIEM to the CounterCraft Deception Director. This also generates a CounterCraft notification.

3 Detecting Enumeration of AD Credentials in Shared Resources

A series of shared resources are presented that are seeded with additional breadcrumbs. The breadcrumbs lead the attacker to CounterCraft Deception Hosts within the **deception environment.** The shared resources will point away from the **production environment** entirely with the aim of increasing the dwell time of the attacker in the **deception environment**, all the while gathering more detailed data on their tools, techniques, procedures and more importantly, their motivation.

So far, we've talked about the basic concepts behind a typical campaign deployed to protect an Active Directory installation using the CounterCraft Cyber Deception Platform. To recap, the deception technology is used to detect this activity in three distinct areas. Next we'll focus on addressing the infrastructure required to achieve the steps already described.

Each area is part of a carefully structured Deception Campaign deployed from the CounterCraft Deception Director. So let's look in closer detail at the infrastructure required to support the Deception Campaign.

To Detect Enumeration of Active Directory Credentials on Production Endpoints:

Breadcrumbs seeded on Endpoint Workstations

To detect malicious activity on production endpoints, carefully crafted breadcrumbs are distributed to the endpoint workstations. The distribution is carried out using the CounterCraft BC-CLI tool, which is designed for massive distribution of breadcrumbs across an enterprise network. The breadcrumbs deployed at the endpoint workstations provide the following misinformation, linked to the concepts described previously:

- ✓ **Location and credentials** for deception hosts.
- ✓ **User credentials** that point to a Deception AD Domain Controller.
- ✓ **PowerShell scripts** with information that leads to either the deception AD Domain Controller or the shared resources located on the Deception Hosts.
- ✓ **Auto login credentials** for the shared services hosted on the deception servers mentioned below.

To Detect Enumeration of Active Directory Credentials on a Production Active Directory Domain Controller:

Breadcrumbs seeded on the Production AD Domain Controller(s)

As mentioned before, to detect malicious behaviour on the Production AD DC (and as with the endpoint workstations) it is seeded with breadcrumbs. The breadcrumbs used in the **production environment** do not in any way affect the function or behaviour of the domain controller. The following examples show a broad spread of the types of breadcrumbs typically used:

- ✓ **Fake users** – The fake users are not linked to any real services, but information pointing to additional deception services is included in the description field.
- ✓ **A fake GPO** – The fake GPO points to a Deception SYSVOL containing additional breadcrumbs, such as PowerShell scripts with data that in turn lead to additional deception services located within the **deception environment**.
- ✓ **Fake resources** – The fake resources are not linked to any real resource, but as with the false users the description field will contain information pointing to additional deception services.

Typically, Active Directory logs are sent to a SIEM. So, to log activity from the breadcrumbs on the Production AD DC, the CounterCraft Deception Director takes in a feed from the SIEM. This allows the Deception Director to monitor breadcrumb activity and provide fully analytical coverage for the Production AD DC without having to install the CounterCraft Agent. Alternatively, a direct log feed from the AD DC can be integrated into the Deception Director.

To Detect Enumeration of Active Directory Credentials on a Production Active Directory Domain Controller:

A Deception AD Domain Controller

To support the Deception Campaign, at least one Deception AD Domain Controller is deployed. This obviously requires a Windows server. The server is fully instrumented by installing the CounterCraft Agent that reports all activity to the CounterCraft Deception Director. The installation and configuration of the Agent is carried out from the console.

Specific Active Directory Event-Types have been created to flag AD specific activity and create the appropriate Notification rules within the Deception Director. The number of deception AD Domain Controllers to be deployed depends on the network topology to be replicated, but the recommendation is to deploy at least one Deception Host Domain Controller per domain or group of domains (known as a 'forest').

To Detect Enumeration of AD Credentials in Shared Resources:

Shared Resource Deception Hosts

The Deception AD Domain Controller, the false credentials and some of the other breadcrumbs will point to shared services running on a series of deception hosts. These deception servers are fully instrumented, high-interaction honeypots. The deception story behind these servers is typically that they are new or test systems under evaluation in a development environment.

Examples include:

- **An IDS server** – An Intrusion Detection System, acting as a double bluff security system. It is fully instrumented and will not only report as a high-interaction honeypot to the Deception Director but also function as a fully operational IDS – reporting to the SIEM – to lend credibility to the deception scenario. Any malicious activity on this device will be captured by the CounterCraft Agent.

- **File Shares** – A Windows file share server, loaded with populated folders. Each folder will have an SCF file that reports to the Deception Director if the folder is touched. Additionally, beaconing documents are distributed within the folders that will also report to the Deception Director if opened. This not only allows access to the beaconing documents to be logged, but also to map the path of the attacker as they navigate through the folders.

- **A Web Portal** – A web server and application. Typically, this is set up as a copy of an internal portal site. It can be useful to capture any credentials used in login attempts. All interaction with the web application (and server) is captured and sent to the Deception Director. This can either be a Windows or Linux server depending on what is more credible given the **production environment.**

- **A Database Server**: A database server in development, or set up to provide support to an application under development. The database is set-up and populated with fake data. Typically, this is a Windows server with the most appropriate database to the Production infrastructure.

- **Clone(s) of Production Server(s):** In more advanced deployments this is an option for the AD campaign. It involves placing a cloned copy of an existing production application server into the **deception environment**, fully instrumented with the CounterCraft Agent. This allows any malicious activity to be instantly detected, and provides a valuable intelligence on exactly how an attacker would interact with a real production system.

To conclude, the use of deception technology to protect an Active Directory installation requires the use of a carefully structured Deception Campaign, with some or all of the components described here.

Any deployment of this sort would be carried out in full cooperation with CounterCraft or one of our trusted partners. It goes without saying that AD is such a critical component, that protecting it correctly is key. Using deception technology provides a way to add an extra layer of protection to your existing security systems, to detect, investigate and control any potential attack.

Next Steps

Find out more by getting in touch. We are only too happy to explain what we do and how we can help you get the best out of deploying deception – from an initial conversation or simple demo, to a fully featured deployment.

References: [1] DESMOND, Brian; RICHARDS, Joe; ALLEN, Robbie & LOWE-NORRIS, Alaistair G. "ACTIVE DIRECTORY" - 5th Ed. O'Reilly, 2013

Learn more about CounterCraft Deception

Download our latest documents at

countercraftsec.com

or if you prefer contact us at

craft@countercraftsec.com