THE **TOP 20** LATERAL MOVEMENT TACTICS

A comprehensive look at how hackers

- Find assets
- Escalate privileges
- Move around your network

SMOKESCREEN

Table of Contents

Introduction

Just a few short years ago, a sophisticated data breach was newsworthy, not just because of the damage done, but because of how rare it was for a technically adept hack to occur. No longer. Advanced cyber threats have come of age.

The modus operandi, or ‘kill chain’ of the modern hacker has matured. An quick entry through spear phishing, web application flaws, or misconfigurations, followed by a systematic compromise of multiple assets, and finally, achieving the mission goal.

The ability of a skilled attacker to bypass traditional protection mechanisms means that cybersecurity is now more about the ability to detect and react rapidly, than it is to try and close every possible loop-hole.

Contrary to what most may assume, the most critical phase of an attack is not the initial exploit, or the final data exfiltration. It’s the middle phase, where the attacker seeks out assets, gains additional privileges and silently moves from system to system, subnet to subnet, closer to his final goal.

This phase is called **lateral movement**, and it is the place where the attacker spends the most time, and is also the most vulnerable to detection.

The Kill Chain

Reconnaisance
Email / web / USB

Delivery
Email / web / USB

Exploitation
Social engineering / misconfigurations / exploits

Persistence
Maintaining multiple points of access.

Command & Control
Remote manipulation of the victim environment.

Lateral Movement
Asset identification, privilege escalation.

Mission Goal
Data exfiltration, destruction or modification.

Through years of experience responding to sophisticated breaches, and an extensive analysis of the techniques used in high-profile attacks, we’ve extracted the most common patterns of lateral movement. Each of these is explained, and then we present solutions to detecting the presence of an attacker within the network.

With a better understanding of how attackers move around your network and stay undetected, you’ll be better positioned to find evil in your environment, and better placed to thwart advanced attacks.

Why Catching Lateral Movement Is Your Biggest Win

Conventional wisdom states that ‘prevention is better than cure’. Unfortunately, the attack surface of modern companies is so large, that protection is akin to building a fence around a national border. You can try, but it’s not going to keep a determined attacker out.

Our research shows that 80% of an attack is spent during lateral movement. The actual breach occurs fairly rapidly, and the final goal is quickly accomplished as well. It’s moving from initial breach to the final goal that takes hackers time and resources.

Even the most savvy attacker is operating ‘blind’ once in the network. They may know where the assets are, but they have to move slowly and stealthily to get there.

If you can catch them during this process, it’s game over for the attacker.

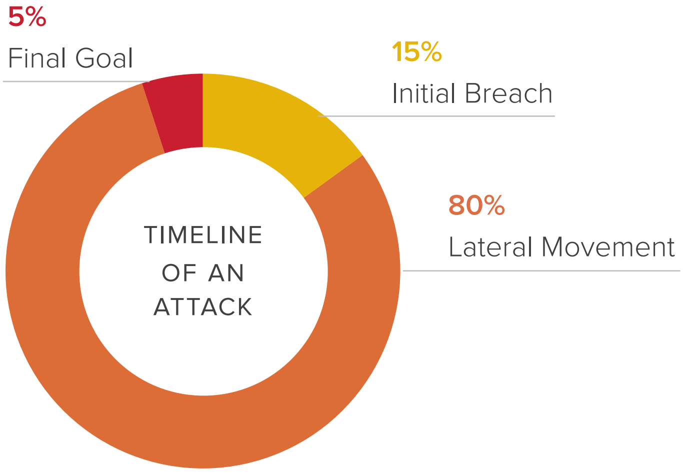

It takes 7 months
before breaches are discovered.

Under 4% of alerts
are even investigated.

The Challenges

On average it takes over 7 months from initial compromise to a breach being discovered. So if detecting lateral movement is so powerful, why aren’t more companies doing it? It’s not for want of trying.

Unfortunately, monitoring internal networks is hard. Companies have tried log analysis, SIEM’s, anomaly based detection and machine learning. But the volume of data is in petabytes, and even the best predictive analytics solutions generate a huge number of false positives.

The problem is so bad, that in the average company, less than 4% of alerts are even investigated! This is because the volume and irrelevance of alerts leads security teams to disable or ignore these monitoring solutions.

Lateral Movement - Techniques, Tactics & Procedures

So how does an attacker move laterally on the network? The tools may change, but the basic strategy remains the same — Gain access to a lower protected, lower privileged asset, escalate privileges, and then start seeking out interesting targets on the network.

After studying the specifics behind numerous breaches, these are some of the most common lateral movement techniques that we have observed (not in order of prevalence).

Psexec

Developed by Sysinternals prior to their acquisition by Microsoft, psexec and the entire pstools suite lets administrators to remotely control Windows systems from the terminal.

Attackers love psexec for its ability to upload, execute and interact with an executable on a remote host. Since it works from a command line, it is easily scriptable, and doesn't alert the remote user to its operation.

Since its also a legitimate system administration tool, it is invariably not blacklisted or detected by antivirus solutions. You will observe this common pattern of using legitimate tools to avoid detection.

File shares

We have yet to see an internal network environment that doesn't have half of what an attacker is looking for available over Windows file shares.

As the mainstay collaboration mechanism, file shares are used both on central file servers, as well as by individual users. They often contain customer databases, details of additional systems, operating procedures, and useful software.

The built-in administrative shares for the full hard-drive are also extremely useful to attackers with elevated privileges.

Once again, these are legitimate traffic channels that go unpoliced by monitoring solutions.

Remote desktop

What better way to control a victim asset than full interactive GUI access, using a legitimate control channel that is built in to almost every version of Windows?

Once attackers have valid credentials, terminal services, or RDP as it is commonly known, is the weapon of choice to gain interactive access to the assets. Since RDP sessions are encrypted, they're opaque to monitoring solutions (which in any case would not flag them as they are such a common legitimate administrative mechanism).

Powershell

As sandboxing technology began catching malware without signatures, attackers moved to 'living off the land', or avoiding malware in their attacks and using built-in operating system capabilities to replicate malware functionality.

The numero uno mechanism is Powershell, Microsoft's object-oriented scripting facility is built-in to every modern version of Windows, and is extremely powerful — attackers have used it to steal in-memory credentials, modify system configuration and automate movement from one system to the next.

Once again, it's legit, so it doesnt' get caught. You're noticing the pattern aren't you?

Port-scans

The simple port-scan — possibly the only technique that has remained virtually unchanged from the days when hacking was a romantic pursuit by curious techies.
The portscan is used to quickly identify services of interest — typically web applications, database servers and remote access functionality.

While a full-blown port-scan is easily detected, "low and slow" scans get past practically any network monitoring system.

Despite Nmap's plethora of features, attackers don't need to bother with bells and whistles like OS detection, or script scanning. Just simple TCP connects are sufficient for finding targets.

WMI

The Windows Management Instrumentation framework is Microsoft's built-in system to manage the configuration of Windows systems. Some consider it SNMP on steroids.

WMI can be used to start remote processes, query system information, or even, as has recently been demonstrated, store persistent malware that does not touch the disk in a traditional sense.

Attackers make extensive use of WMI as a means of quickly enumerating system information to classify targets.

Scheduled tasks

The simple Windows 'at' command allows an attacker to schedule a task to execute on either a local or remote system.

This ability is not used just for timing execution — in many circumstances, scheduled tasks will run as the SYSTEM user, letting an attacker escalate privileges to complete control of the host on which the scheduled task runs.

It's also a great way to schedule batch jobs that may use CPU or bandwidth, such as zipping up folders and transferring them over the network. By scheduling the task out of office hours, there is less chance of detection.

Token stealing

While a fairly recent technique in the public domain, stealing tokens from memory has become all the rage, and is used in almost every attack these days.

Tools such as mimikatz and Windows Credential Editor can find service accounts in memory, create Kerberos tickets, and elevate an attacker from normal user to domain administrator in a few seconds.

Fully undetectable versions of these tools are easily available, while savvy hackers have implemented the functionality in Powershelgl to avoid detection.

Pass-the-hash

Due to how the NTLM protocol works, attackers can use the encrypted hash of a password to authenticate to remote services without knowing what the plaintext password is.

After obtaining the password hashes, the attacker simply passes them on to other services, without having to undertake dictionary or brute-force attacks on the hash itself.

This technique has been superseded in many attacks by token stealing, but has still been used to devastating effect in recent breaches such as the Target Corporation attack.

Active directory

The Holy Grail of victim discovery and control. Active Directory is the 'phone book' of your network, and practically anything of value has a place in it. It's no surprise that the first thing attackers go for is the list of AD computers.

The names in this list are the starting point for categorising targets. A quick search for databases, backup systems, SCADA controllers, you name it, they read the names.

Furthermore, breaking into the domain controller or gaining domain administrator privileges gives hackers the keys to the kingdom, at least as far as Windows environments are concerned.

Remote registry

Everyone knows that the Windows Registry is the heart of the operating system.

While usually used as part of a larger technique, the ability to remotely manipulate the Windows registry can be used to disable protection mechanisms, remove auto-start programs and services, and install persistence mechanisms.

Admin shares

While we've already spoken about file sharing, the built-in ADMIN$, C$ etc. shares deserve special mention, because they are used in practically every attack we have analysed.

The ADMIN$ share is the mainstay of psexec style attacks, and gives complete access to the %SYSTEMROOT% folder.

Meanwhile the per-partition hidden shares give complete read-write access to the entire hard-drive of the remote system. What more would an attacker want? Once again, virtually undetectable as it has legit uses.

Stolen credentials

When companies invest heavily in anti-malware, they're missing the fact that attackers moved on years ago (despite what sandboxing and antivirus vendors may say).

Gaining legitimate credentials is easy once on the internal network, and once you have them, it's far safer for the attacker to use legitimate administration channels and stolen credentials to complete their mission.

Every single attack we analysed made use of legitimate credentials that had been stolen, phished, cracked, key logged or gained in some other fashion.

Notably, the FIN-4 financial APT group ran a highly successful campaign against Wall Street using nothing but stolen credentials.

Breached host analysis

If you break in to someone's house, you're bound to look around for things of value.

Similarly, when attackers breach an initial host, they pillage it for information that can help them move further. This includes passwords in text files, details of other systems, operating procedures, and even screen captures of how the user is working. They'll even just figure out the internal hierarchy and politics of the organisation to craft social engineering attacks.

Consider this foot printing behind the firewall, and, when properly done, it sets the stage for devastating attacks that can lead to a cascade of compromised systems.

Central admin consoles

The modern attacker is lazy. Rather than spending time breaking into individual hosts, it's far less effort to break in to the system that controls them all in some way, and treat it as a mini legitimate botnet controller.

ATM controllers, point-of-sales management systems, remote management tools like Ansible and Salt are all great targets as they give an attacker a "one shot, thousand kills" capability.

You can consider Active Directory the grand-daddy of central administration systems.

Network sniffing

While switched networks have made promiscuous mode sniffing less of an issue, attackers still gain tremendous value from setting up network sniffing on a high-traffic server to gain access to credentials of customers and other information.

User segments are usually subjected to man-in-the middle attacks like ARP spoofing, explained below.

ARP spoofing

It's been around for ages, but it's still used. Generating gratuitous fake ARP requests and replies can let attackers interject themselves in communications in switched networks in a classic man-in-the-middle attack.

While these attacks have fallen out of favour somewhat, they're still not detected on many networks, and can be extremely damaging if an attacker finds the right hosts to poison.

Email pillaging

A famous Romanian hacker once said that over 35% of the information in a company is stored directly in email.

It's no surprise then, that gaining access to email inboxes, either on the workstations, the server, or through webmail gives an attacker tremendous leverage.

Besides the obvious information gathering and second tier spear phishing attacks, there are instances of attackers watching an incident response process being discussed over email, and modifying their tactics accordingly.

Software deployment

Everybody loves centralised software installation systems. Attackers included.

These are literally systems that legitimately are allowed to execute arbitrary code on thousands of hosts across the environment. While you can consider them a central administrative console, they bear special mentioning with respect to attacks against point-of-sales systems and ATM networks.

There is some information that the Sony attackers made use of Microsoft's built-in software distribution mechanisms to spread on the network as well.

VNC / Ammy Admin / Teamviewer

The venerable VNC GUI remote access is still used in thousands of companies for remote technical support.

Newer solutions like Ammy Admin and Teamviewer offer more capabilities, including file transfers and persistence features.

All of them are extensively used by attackers who target one administrator with legitimate access, and then use those credentials to access all the other systems in the environment with his or her credentials.

If you use a central remote support system, it requires very strong protection and monitoring.

The Solution

Detecting lateral movement is not just possible, it's within reach of most organisations.

Let's start with the key takeaways from our previous study of lateral movement:

- Attackers need to discover assets through scanning, Active Directory or foot printing.
- Using legitimate administrative tools or stolen credentials makes traditional detection very difficult, as monitoring for these will generate thousands of false positives.
- The techniques used keep evolving (for example pass-the-hash to token stealing, sniffing to ARP spoofing), but the basic strategy is the same — Find assets of greater value than the current asset, and compromise them.

Therefore, the ideal solution has to be:

- Agnostic to the attacker's tactics
- Should not generate false positives
- Able to detect movement across assets

We've found that the solution is to focus on what the attackers want, and to make them think they've got it. The solution is deception.

We create virtual "decoys" that perfectly mimic the assets attackers want, and we place them where an attacker has to find them during lateral movement. When they find them, they ignore real assets, waste their efforts, and reveal their presence.

Our ILLUSION**BLACK** system automates the deception and watches for lateral movement 24/7, without generating false positives.

ILLUSION**BLACK**

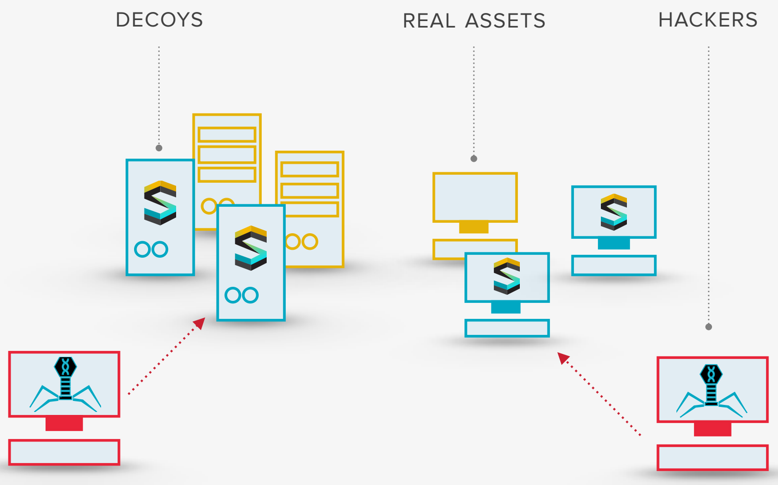

About Smokescreen

Smokescreen Technologies Pvt. Ltd. creates the next generation of specialised cyber-security systems.

Our proprietary ILLUSION**BLACK** platform defeats hackers in a manner that is free of false alerts, and requires no changes to existing networks. It effectively solves the problem of multiple avenues of attack and the limited response capabilities that companies have.

Our ‘active defence’ philosophy is the result of decades of experience securing the most highly targeted organisations globally, and has proved its effectiveness time and again.

For a free consultation, contact us at:
www.smokescreen.io | info@smokescreen.io