External Attacks & Privileged Accounts

5 STEPS TO CONTROL THE THREAT POTENTIAL

BY NICK CAVALANCIA

Table Of Contents

Introduction: The Reality Of The Privileged Threat Potential 1

1. Understanding The Threat Potential: External Attacks & Privileged Accounts 1

2. Responding To The Threat Potential: A Layered Approach 3

3. Envisioning The Threat Potential: Discovering Your Vulnerabilities 6

4. Managing The Threat Potential: Managing Enterprise Passwords & Vulnerabilities 8

5. Monitoring The Threat Potential: Watching Privileged Account Behavior With Threat Analytics 10

Summary: The Privileged Threat Potential Under Control 13

INTRODUCTION: THE REALITY OF THE PRIVILEGED THREAT POTENTIAL

The job of IT used to be focused around implementing and supporting services that kept the business running. But as connectivity has increased, and those with development skills "use their powers" for evil instead of good, IT's focus has shifted to not just *include* security, but to weave it into the very fabric of every part of the network environment.

Everything from email, to files, to printers, to databases, to web applications (and everything in between) needs to be secured because the lone attacker has evolved into sophisticated organizations and now leverages as-a-Service-type automated hacking tools. *If someone wants to attack your organization, it's only a few clicks away to get started.*

Without some kind of definition of who's attacking you and an appropriate response plan, it feels a bit like being out in the middle of the woods jerking at every crack and howl, wondering if that is the sound of an attack you need to be worried about. And, that's no way to address the problem.

So, what kind of threat should you be vigilant against?

External attackers have consistently remained the #1 threat to organizations and remain so as the primary actor in roughly 80% of all attacks. They are primarily interested in financial gain, use a mix of malware and stolen credentials as their key methods of attack, and are laser-focused on both your server and workstation endpoints.

The challenge is that the threat source and methods are constantly changing and a given threat may or may not rear its ugly head within the network "walls" of your organization. That's why you need to focus on the threat *potential.*

To be clear, it's not FUD, because the threat is real. The reason to focus on the *potential* is that the magnitude of the threat you may face really depends on you - how serious you believe external attacks to be credible, what steps you take to see just how vulnerable you really are, and how you choose to respond in the form of implementing layered security.

This white paper focuses on the problems associated with external attacks and the resulting assault on privileged accounts. It provides guidance in how to control the potential threat to your data, applications, and systems by walking you through 5 key steps:

1) Understanding the Threat Potential: External Attacks and Privileged Accounts

2) Responding to the Threat Potential: A Layered Approach

3) Envisioning the Threat Potential: Discovering your Vulnerabilities

4) Managing the Threat Potential: Managing Enterprise Passwords and Vulnerabilities

5) Monitoring the Threat Potential: Watching Privileged Account Behavior and Threat Analytics

1. UNDERSTANDING THE THREAT POTENTIAL: EXTERNAL ATTACKS & PRIVILEGED ACCOUNTS

So, you're under attack, right?

Most of you are probably in the "it won't happen to us" category. But, you're reading this white paper, so there is a part of you that recognizes "well... maybe it could happen to us." And that's good. Admitting that you have a problem is making the first step.

In reality, external attacks are on the rise – most organizations have experienced a targeted attack, including via hacking, malware, and phishing. And these attackers are looking for very specific industry verticals and data sets - healthcare records, credit cards, intellectual property, etc. – all, of which, have real monetarily value.

Long gone are the days of the attacker being some guy in a hoodie sitting over his computer in a dark room. Today's attackers are sophisticated organizations that are, in some cases, real companies complete with salaries, benefits, healthcare, etc. In essence, crime has become a business. And business is booming.

So, how are they going gain access to that valuable data you've got locked up?

Attackers enter an organization by taking advantage of an unpatched system, or via malware or phishing attacks, or stolen credentials and rarely are left in control of just the right server with just the right privileged account. So, they use that initial endpoint as a foothold and begin a process of gaining access to systems and data that aid in their overall goal. For most attackers, the goal is exfiltration of data, while others can be focused on data destruction, or even holding data for ransom.

Whether or not the attacker knows specifically what data or system they're looking for, to get there, they know traversing multiple systems, requiring supplementary sets of credentials along the way, is a necessary step to facilitate additional footholds within the organization.

WHY TARGET PRIVILEGED ACCOUNTS?

You could just stop with the obvious reasoning of "because those accounts have access." But stopping there would be short-sighted. Stopping there makes you think about only those accounts that have direct access to critical, sensitive, or valuable data or similar systems within your environment, narrowing your focus on a small number of accounts.

In reality, an attacker is interested in gaining control over as many privileged accounts as they can possibly get their hands on. Everything from a user with access to uninteresting marketing files on a file server, to the domain admin account, and everything in between. They'll need these accounts along the way as they progress laterally across you network, until they reach the data or system in question, then they need that specific privileged account you were thinking of.

Then how do attackers find and access privileged accounts?

To better understand how privileged accounts are used as part of an external attack, let's take a look at a somewhat typical set of actions an attacker may use once they've gained access to an endpoint.

ATTACKING PRIVILEGED ACCOUNTS

Remember, privileged accounts aren't just those that have admin access to a database. They can be a normal user account with admin-level privileges on their local workstation. Again, attackers will use every credential they can get their hands on as a potential means to achieve lateral movement. Remember, assuming the attacker already has some level of access to an endpoint, they follow a simple, yet effective, process:

1) Identify and obtain credentials with privileged access

2) Access another endpoint

3) Repeat until getting to the desired system or data set

So, how do attackers identify and take advantage of privileged accounts?

There are a number of steps they take:

1) Surveying local user privileges – the initially compromised account may just be a low-level user... or they may be a domain admin. The attacker needs to find out. So they enumerate the groups the local user belongs to (using something as simple as NET.EXE). Additionally, attackers identify users that are in the local Administrators group (and any other group that appears to have some level of elevated privileges), to see exactly which account will elevate their permissions, giving them a greater foothold on that machine.

2) Attempt to logon with elevated privileges – an obvious choice, given the local admin account may have no password in some cases. Trying out a simple *Run As* would do the trick.

3) Bait and Wait – If no privileged access is uncovered, the attacker is patient and will wait until they have the right set of credentials. They place an infected file – such as a Word document with a malicious macro causing a buffer overload – up on a common folder in a file server via a mapped drive, waiting for another user to open it and infect a second endpoint.

4) Use tools to crawl through the endpoint's memory – tools like *mimikatz* can be used to obtain plaintexts passwords, domain credentials, NTLM hashes, and Kerberos tickets, potentially of privileged accounts left in memory. These credentials can be remnants of connections to other systems, someone from IT logging on to support the user, etc. Each of these sets of credentials can be utilized via additional tools using attack methods such as *pass-the-hash, brute force,* and *pass-the-ticket.*

5) Move laterally within the organization – if a set of credentials is found to provide access beyond the initially infected machine, attackers will work to connect to and infect additional systems using methods like WMI, remote desktop sessions, PowerShell remoting, and SMB file replacement.

These actions flow in a cycle, repeating until the attacker finds what they seek.

UNDERSTANDING THE THREAT TO PRIVILEGED ACCOUNTS

It's evident to both you and attackers that privileged accounts are the only way to gain access to valuable data within your organization. Common tools and methods are used to both identify and use these credentials, making the work to access and retrieve information more of a process than a challenge. All this adds up to the simple statement – the threat is real and credible.

So, what's the appropriate response to the threat?

In the next chapter, we'll take high-level look at how you should respond to the threat potential, the types of solutions available, and the features and functionality necessary to work towards eliminating the threat potential. We'll also provide some practical advice on steps you can take to respond immediately to the threat.

2. RESPONDING TO THE THREAT POTENTIAL: A LAYERED APPROACH

In the last chapter, we covered some ways external attackers attempt to gain access to privileged credentials and the reasons why. And because credentials are the number one attack vector of external attackers, we want to focus on what you need to do specifically to protect your privileged accounts in response to the growing threat of external attacks.

You can think of an attack in three layers – each one representing a higher degree of threat potential to your organization:

Access – Identifying and gaining access to credentials that have elevated privileges via enumeration of groups, testing of local accounts, and use of tools to pull credentials from memory is necessary to expand an attacker's presence beyond just the initial foothold.

Entry – Control over an endpoint is established and used as a foothold. Phishing, malware-laden websites, and exploiting known vulnerabilities are the most common methods used to gain entry.

Activity – Putting the found privileged credentials to good use, attackers attempt to move laterally across your network by accessing additional systems via methods such as RDP, WMI, and PowerShell.

Remember, these three layers occur in every attack, which makes the job of IT a little easier in identifying how to put up a matching layered defense. But also consider that the specific types of malware used and actions taken can vary and evolve, making this an ever-changing threat.

Your response to the potential threat of an external attack needs to exist at all three attack layers – that is, before an attacker gains any kind of entry into your organization, before they can identify and utilize credentials, and as they attempt to utilize them. In all cases, the responses are largely too complex to be accomplished manually and, therefore, will require you to put solutions in place.

So, what are the appropriate responses for each attack layer?

ENTRY RESPONSES

These responses are designed to proactively stop an attacker before they gain entry. There are two basic means of entry – malware and known vulnerabilities. To protect against each, your response needs to include the following kinds of measures and solutions with the goal of stopping an attacker at the front door as often as possible:

1) User Education – While it's going to be tough to keep users up-to-date on every type of phishing and spear phishing attack, educating them on what to look for and what to do should they find a suspicious email (e.g. notify IT via a specific email address) is a solid first step.

2) Endpoint Security – It's not the phishing emails themselves that are dangerous; it's the malicious attachments that shouldn't be accessed by users. Endpoint protection solutions (AV, privileged access management, EDR, etc.) can provide a layered defense, keeping malicious code from running.

3) Vulnerability Management – While phishing and malware are ever-evolving threats, in vulnerability-based attacks, the overwhelming majority of exploited vulnerabilities occurred more than a year after the vulnerability was identified and published as a CVE! This makes vulnerability scanning and patching of your operating systems, enterprise applications, and known "applications of access", such as Flash and Java, a critical part of your proactive response. This is one response you can use native tools to address, but those of you in enterprise environments will likely lean toward a solution that automatically identifies and remediates vulnerabilities across the entire network.

These responses provide a defense between the attacker and your endpoints. But, should an attacker gain entry with a new, yet unseen, variant of attack or unpublished vulnerability, you're going to need some additional security measures to specifically address the *Access* and *Activity* layers of an attack.

ACCESS RESPONSES

You run a greater risk should you simply assume that, by having vulnerability management and antivirus/antimalware, you are completely protected. Once an attacker gets in undetected, it is truly only a matter of time before they gain privileged access to parts or all of your environment. So, you need measures in place that prevent access to privileged accounts in the first place, including:

1) Least Privilege – By ensuring users requiring elevated permissions have separate accounts for low-level tasks (e.g. email, web surfing, etc.) and privileged tasks, you lengthen the time it takes for an attacker to obtain privileged credentials, increasing your chances of detecting them. Reducing privileges, such as via a privileged access management (PAM) solution, also helps eliminate lateral pathways.

2) Privileged Password Management – Beyond just the centralized secure storage of passwords in something other than an Excel spreadsheet or Word document, there are a few ways PAM helps to secure privileged accounts that keep attackers at bay. One is the concept of password rotation. By having password values rotate after each use, any credentials or credential remnants left in memory become useless. Another use of PAM is proxied access where, rather than providing the password directly to the user (which, then can remain in memory for later use by an attacker), PAM solutions can put a user directly into a server-specific session, eliminating local storage or typing of a password, lowering the risk of misuse of that account. Depending on the PAM solution, there are many other possible features that will help to prevent access to privileged accounts.

In many ways, PAM is proactive in nature – you can restrict, obfuscate, and control access to and use of privileged credentials. But it is still possible for an external attacker to utilize a keylogger and gain access to a privileged session. This is why you need to be concerned with the Activity layer of an attack.

ACTIVITY RESPONSES

The mere use of a privileged credential may not look suspicious – an attacker can use the very same credential you do, from your machine, to the same server you do. So, it's important for your response to include an ability to detect inappropriate use of privileged accounts. Activity responses should include:

1) Behavior Analytics – Here you're looking for anomalies in the when, from where, to where, and how accounts are used. A logon to the database server after hours or on a weekend, from an endpoint that's never tried to access the database server? Yeah... that's suspect. The trick here is you need to establish what normal looks like – and from many, many angles. And then determine when something looks out of the ordinary – which equates to detecting the inappropriate use of a privileged account.

2) Session Recording – As a last-ditch effort should an attacker gain entry, access some credentials, and use them on another endpoint, you need to determine how those credentials were used. Did data get exfiltrated? Was malware inserted into a server? Which databases were compromised? You need to know. Session recording would allow you to replay the session and identify specifically what actions were taken.

PUTTING A RESPONSE IN PLACE

You can't wait for an attack to occur before you respond - you probably won't know it's happening anyway. Thinking about the threat potential in layers provides your organization with an ability to respond in kind, putting a stop to malicious actions when and where they occur. Responses will come in form of both solutions and IT activity – you'll need to decide which methods of response best fit your security strategy, resources, and budget.

In these first two chapters, we've provided some high-level detail on what the privilege-related threat looks like and what, generally, you need to do to address the threat potential. In the remaining three chapters of this white paper, we'll cover specifically how to identify, manage, and monitor the threat potential in your environment.

3. ENVISIONING THE THREAT POTENTIAL: DISCOVERING YOUR VULNERABILITIES

In the prior chapter, we outlined a number of responses appropriate for each layer of an attack. These responses involved both solutions and putting proactive security measures in place, creating a layered defense against the threat potential. But there's still a problem – where should you be placing your efforts and focus? Should it just be on your servers? Do only your Domain admin accounts need to be managed?

In reality, the responses in Chapter 2 are definitely what you should do, but there's still the question of where you should do it. Envisioning the threat potential involves you understanding exactly where you are leaving yourself vulnerable in terms of both the endpoints that may end up acting as footholds, and the privileged credentials that may be hijacked to give an attacker the access they need.

So, how do you discover what vulnerabilities exist?

First, let's level-set what should be meant by the term vulnerability. As an industry buzzword, the term usually is associated with just patching applications and operating systems. Given the threat potential involves the actual gaining of entry to an endpoint and access to credentials, to truly have a comprehensive understand of where you're vulnerable, you need to expand the definition beyond just patching. Your vulnerability thinking needs to include visibility into every asset you manage and the privileged credentials you don't.

Next, let's break vulnerabilities down into the three aforementioned categories (assets, patching, and credentials), and look at how to best envision where you are vulnerable.

IDENTIFY EVERY ENDPOINT ASSET

Remember, every endpoint is susceptible to an attack. Discovering every client and server system that exists on your network – whether local, remote, in the cloud, mobile, physical, or virtual– is a critical first step. To be clear, we're nowhere near the discussion of published common vulnerabilities and exposures (CVEs) for both applications and operating systems yet; we're simply talking about knowing all the hardware and software assets that exist. After all, you can't patch what you don't know about.

It's critical to be inclusive when it comes to outlier operating systems, applications, and even devices. The goal here is to analyze your risk by first gathering every means of entry into your network. So, that means you look beyond just Windows, and think Unix, macOS, mobile device OSes (iOS, Android, Blackberry) - anything and everything with access to your network. While there may not be much you can do when it comes to known vulnerabilities on some of these devices, knowing they exist will still give you a remediation step in the event of an attack in the form of disconnecting the endpoint from the network.

Once you know what assets you have, it's time to make them as secure as possible.

ASCERTAIN EVERY PATCHABLE VULNERABILITY

Nearly every OS and applications vendor has some kind of automatic update built into their product, but envisioning the threat potential is so much more than just patching. It's really about understanding, from a patching perspective, where you're vulnerable and whether you've filled in those known gaps in security.

The impetus for patching should be squarely based on the need for security, with "keeping up-to-date" as a secondary benefit.

Year-after-year, it's evident from attack and breach data that organizations are either not patching or they're not patching the right things. Often organizations don't have a full inventory of the assets they have that need to be patched (which is why the first step in this chapter is so critical), and/or they have no means to verify and ensure the assets all fully patched.

Given you're serious about the threat potential, it may be time to look at an enterprise-class solution that centrally scans, assesses, identifies, and reports on vulnerabilities across as many of your assets (both hardware and software) as possible. So, if you're simply using something like WSUS, you may be missing out understanding which third-party applications and non-Windows operating systems leave you vulnerable, both today and in the future.

To be truly prepared for the threat potential, you need to have a holistic view into the gaps in security that exist both now and as the threat landscape changes. When you are certain you have complete visibility into patch-related vulnerabilities, you should still assume new vulnerabilities would be discovered and taken advantage of by attackers. So, you'll need to also work to understand the threat potential that exists in the form of your privileged accounts that will be the next layer of an attack to be compromised.

UNCOVER EVERY PRIVILEGED ACCOUNT

Your privileged accounts don't consist of just the Administrator account in Active Directory; you have accounts that provide access to data, applications, and systems – all resources an external attacker would like to get their hands on. Because the attacker wants to leverage every account they can compromise to aid in the endgame of accessing protected data somewhere within your network – from a local admin on a workstation, to an Enterprise Admin in AD, and everything in between – you need to know what privileged accounts exist within your network. Doing so will help you envision what that threat potential looks like – and provide you context of what needs to be done next to protect those accounts.

So, what accounts should be included?

As a general rule, any account with more privileges than a low-level user should be part of your discovery. Accounts can exist across every OS in your environment – Windows, Unix/Linux, macOS, etc. and reside within local databases as well as in directory services like AD. You need to include accounts with privileged rights to workstations and servers, service and daemon accounts, accounts with application-specific privileges, those with access to networking hardware, and those embedded within scripts and applications. Additionally, treat SSH Keys as privileged accounts, as they can be used to provide elevated access to one or more Unix servers.

Keep in mind the list you generate can change at a moment's notice, but it's an important first step to at least have some level of understanding where your privileged accounts are so that you can take the next step to protect them... putting them under management.

ENVISIONING WHERE YOU'RE VULNERABLE

You can't protect your environment if you're not aware of the means by which an attacker will attempt to gain entry and access. By gaining visibility into the endpoints you have, the application and OS vulnerabilities that need to be patched, and the privileged accounts that will be the target of an attack, you empower IT to have a firm grasp on the threat potential and the corresponding risk that exists if you do nothing to put these elements under management. In the next chapter, we'll discuss how to put controls around each of these elements, helping to manage and, in some cases, eliminate the threat potential.

4. MANAGING THE THREAT POTENTIAL: MANAGING ENTERPRISE PASSWORDS & VULNERABILITIES

In the last chapter, we discussed the necessity of understanding exactly where in your environment you are vulnerable from the perspective of how attackers attempt to perform the first two of the three attack layers mentioned in Chapter 2 – Entry and Access. (We'll address Activity in the next chapter.)

Entry is often accomplished by compromising an endpoint using known vulnerabilities that have existed, in some cases, for many years. And access is achieved by discovering and utilizing privileged accounts that can provide entry to either the target system or one that gets an attacker closer to their target.

As mentioned in the last chapter, once you have complete visibility into your endpoint assets and patch status, as well as of every privileged account used – regardless of OS, application, or method employed – only then can you truly understand the threat potential as it exists within your environment, and begin to take action to manage it.

So, what's involved in managing the threat potential?

Next, let's break vulnerabilities down into the three aforementioned categories (assets, patching, and credentials), and look at how to best envision where you are vulnerable.

The goal, simply put, is to shut down an attacker's ability to use your systems and accounts against you. To achieve this goal, you need to put proactive measures in place that monitor, manage, and measure the very same elements an attacker would use to ensure those elements either cannot be misused or are not being misused.

Some of you may be thinking, "Ok – so I patch everything and change my privileged account passwords. Got it." But it's not as simple as that. Assuming you've completed the steps outlined in the last chapter and have a solid grasp on just how vulnerable you are, the management necessary needs to ensure that those vulnerabilities aren't just addressed for today, but are permanently addressed with ongoing management in place.

For example, if you decided to change all the passwords today for accounts with passwords older than a year, you'd only be securing those accounts for a very short duration of time. Remember the attacks on an endpoint's memory back in Chapter 1? If an attack occurs tomorrow (after the password reset) and the new password is in memory, the attacker will get control of those credentials.

The management of the threat potential involves implementing a continuous state of protection against both entry and access. Let's look at how you manage against each.

ENTRY PROTECTION: VULNERABILITY MANAGEMENT

A CONSTANT STATE OF VIGILANCE

You'll notice this section is not titled "put patch management in place". There's a reason for that. Making certain every endpoint OS and application installed is patched is part of the answer here. But vulnerability management isn't just about update – it is also about a realization that the definition of where your endpoints and applications are vulnerable is *a daily moving target*.

Regardless of the tactics used, most breaches begin with an attacker exploiting a single external vulnerability on a low-level system. Once inside, adversaries inconspicuously troll around your network, as employees or contractors – capitalizing on privileges to gain access to critical systems and sensitive information. That's why vulnerability management is as much about monitoring the state of your endpoint's patching as is it monitoring them for new types of attacks.

To address the threat, you need more than just patching. After all, nearly every operating system and major application these days has self-updating. That's not the challenging part. Knowing whether every endpoint, every OS, every application, and every device is patched is the challenge. Remember, patch downloads get interrupted, and patching sometimes fails. To be impactful, IT needs to be able to centrally monitor and manage vulnerabilities and the associated patching across as much of your environment as possible.

Proper vulnerability management should have a few capabilities to take that very long list of endpoints, operating systems, and applications, and put some context around both them and your current state of security. Here are just a few examples that demonstrate how gaining intelligence around vulnerabilities and privilege are clearly better together:

1) **Assessing Vulnerability** – being able to look at your endpoints from the malware's perspective – rather than just a list of what patches are available – will help determine how much risk a given endpoint presents to the organization.

2) **Asset Profiling** – similarly-used endpoints should be equally secure. So, knowing All the endpoints in Finance or all the servers housing intellectual property will help in both securing them and prioritizing your efforts.

3) **Patching of Everything** – it's already been stated, but is so important it deserves another go – centrally patch every OS, every application, etc. regardless of manufacturer.

4) **Solid Reporting** – Vulnerability management is all about taking action to secure your environment. A report is only as good as the action you can take from it, so reports need to be contextual, informative, insightful, and actionable.

Having a vulnerability management solution in place will give you both the needed visibility mentioned in the previous chapter, as well as actionable intelligence to best protect your organization, achieve compliance, and communicate risk enterprise-wide.

Using vulnerability data to inform privilege establishes a more intelligent enterprise password security framework - enabling organizations to further tighten least privilege access, improve threat visibility, and, ultimately, make smarter security decisions.

ACCESS PROTECTION: PRIVILEGED ACCOUNT MANAGEMENT

DEFINED & CONTROLLED ACCESS

Attackers have time and tools on their side. Given enough time, they'll eventually obtain a set of privileged credentials. The first step in thwarting the use of these privileged accounts is to put all of them under management, storing every password in a secure database, requiring the checking in and out of the password. Second, is to establish password rotation after each use. These two features alone

would make any use of tools like mimikatz useless, because the credential details obtained will already be out of date.

Additional layers of management include policies around:

1) Limiting the Scope of Use – restricting when credentials can be used, on which systems, from which systems, and by whom all further chips away at an attacker's ability to utilize a set of credentials.

2) Peer-Approvals – requiring a peer or manager to be notified and requesting approval for the use of an account puts a human layer of accountability and security in place. Should you get an email at 3 a.m. on a Thursday night that Bob wants to use the SQL Server admin account, you'd know it's a red flag.

3) Access Methods – Attackers like to connect to remote machines using a variety of methods. Limited access to, say, local logons only, helps to impede an attacker's utilization of any credentials. Taking it a step further, some privileged account management solutions employ built-in establishing of sessions, requiring the use of a management solution to gain access.

Privileged Account Management can be augmented with Vulnerability Management to eliminate threats from unidentified privileged accounts by using the best discovery capability available – a vulnerability scanner – to discover and profile all users, assets and services.

By combining efforts, organizations can share information about assets, users and associated applications, so that IT and Security Ops can make least privilege and security decisions based on their collective information, working together, and not have to settle for using fragmented parts.

THE THREAT POTENTIAL: MANAGED

With a solid understanding of everywhere you're vulnerable, there's no reason that you can't manage the threat potential. The keys are to bring all vulnerabilities under management, and to have that management solution constantly update the environment to ensure the highest level of security.

By putting vulnerabilities – of both the patching and credential kind – under management, IT increases its visibility into where the threat potential can exist, and take proactive measures to ensure those threats never come to fruition.

In the next and final chapter, we'll look at the third attack layer – Activity – and explore the steps to properly monitor the usage of privileged credentials to identify improper use.

5. MONITORING THE THREAT POTENTIAL: WATCHING PRIVILEGED ACCOUNT BEHAVIOR WITH THREAT ANALYTICS

In the previous chapters, we've outlined the need to understand what privileged account threats you're facing, where the threat potential exists, and how to protect yourself against it by managing both vulnerabilities and privileged account use. This leads up to an assumption that you've eliminated the threat potential, right? But this thinking is akin to locking all doors to your home during the zombie apocalypse and assuming no zombie will ever get in. If you were truly concerned about security, you'd be still walking around your home's interior making sure no zombies have made it in via some entry point in your home you didn't notice, secure, or plan for.

So, it's important that you watch the use of your privileged accounts as a way of identifying if and when an external threat has entered your network, compromised an endpoint, and now controls a privileged account.

Back in Chapter 2, we simplified an external attack down to 3 layers – Entry, Access, and Activity. In the prior chapter, we addressed the first two layers through implementing a combination of vulnerability and privileged access management.

In this chapter, we want to address the threat potential by watching for any unusual or inappropriate privileged account activity as a final level of protection using two methods of detection: monitoring and analytics.

MONITORING PRIVILEGED ACCOUNT ACTIVITY

Up until this point, every bit of protective response has been about keeping systems and accounts out of the hands of attackers. But you need to assume, at some point, an attacker will gain access to an endpoint and a privileged account. So, monitoring the actions taken with privileged credentials is a critical step in providing comprehensive protection against your threat potential.

So, what should you be monitoring?

Remember, we're focused here on the activity of an attacker. So, specifically monitoring activity can be broken down into a few areas:

1) **Check-In/Out of Credentials** – Assuming you have some level of Privileged Account Management in place that requires the checking in and out of credentials, along with an ability to establish and restrict session access, monitoring the details of an account's usage can provide valuable insight into whether the usage will be appropriate or not.

2) **Use of Credentials** – More frequent use of credentials can indicate misuse of that account.

3) **Session Details** – By monitoring privileged session details (e.g. when, from where, on which system, etc.), IT can quickly identify unusual usage of accounts. For example, if an account is used to log onto a server not normally accessed by a given account, it's a leading indicator that there may be foul play.

All of this monitoring is somewhat predicated on the basis that you'll know the difference between appropriate and inappropriate access using some very noticeable changes in usage.

But what if the changes are subtle? That's where the right analytics provides a layer of protection.

IDENTIFYING THREATS WITH ANALYTICS

In a perfect world, someone would use an account, IT would be notified, and the anomaly would be instantly spotted. But let's say an attacker gains control of an account that has access to many systems and is simply logging onto a system that has been previously accessed by that very same account. Nothing looks entirely out of the ordinary, and yet, you still need to be able to tell whether that access behavior is good or bad. This is where threat analytics comes into the picture.

Threat analytics uses behavior baselines for a given privileged account, empowering it to spot anomalies in behavior that people would simply miss, such as:

1) **Frequency** – How often is a given account checked out? Used? Used on a specific endpoint?

2) **Duration** – Is a given privileged session longer than normal?

3) **Date/Time** – Is the privileged account being used after hours? On weekends? Or just way out of the normal scope of regular use?

4) **System Changes Made** – A key indicator of misuse – attackers install malware to maintain control and obfuscate their presence.

Threat analytics looks at all these factors in an effort to discern shifts that, together, may indicate a potential threat.

Take, for example, the mix of logging onto a server normally used by a given account but used at a time that doesn't fit normal behavior. It could be an obvious one – like 3am on a Saturday, which IT would likely pick up on their own. But it also could be a time of day during which the specific account has never before logged onto that specific system. The job of analytics is to track what's normal and notify IT when something is out of the ordinary.

ANALYTIC REPORTS IT EXECS WILL LOVE

Analytics can be equated with tons of data having varying levels of relevance to the IT organization. So, what kinds of analytics reports would your CIO want to see? Here are a few examples:

1) **Top 10 most vulnerable systems** – this kind of report provides direction on which systems need remediation and monitoring efforts.

2) **Password Age** – old, unused passwords can indicate inactive accounts that could provide backdoor access.

3) **Scheduled Password Changes** – this type of report should provide a birds-eye view of which accounts' passwords are going to be expiring soon, and the password policy dictating how complex passwords will be.

RESPONSIVE MONITORING

All of this monitoring and analytics is great, but without an ability to do something about it, it's somewhat useless. So, solutions should also include the ability to respond to inappropriate use through IT actions, such as pausing access to an active session, killing that same session entirely, and temporarily removing access of that specific account to that endpoint.

Let's take inappropriate access a bit further. Let's say that an external attacker goes completely unnoticed, accesses a server, performs some dastardly deed (like stealing data), and logs out. Your monitoring needs to provide some level of reactive detail so that once you do realize you've been breached you can assess the extent of the breach.

Monitoring also needs to include an ability to generate an audit trail of usage by recording privileged sessions for replay later. These replayed sessions provide complete context around the actions taken by watching the session, like using a DVR. You shouldn't be limited to just knowing a privileged account was used, but also have monitoring in place that informs you what was actually done with it.

SPOTTING A THREAT WHEN (NOT IF) IT HAPPENS

While we'd all like to think the initial layers of protection will keep the attackers out, you need to plan for the possibility that, as they get smarter, they will get in. Monitoring and Analytics play a critical role as an ever-present watchful eye over the usage of your privileged accounts, looking for misuse that may indicate a threat.

SUMMARY: THE PRIVILEGED THREAT POTENTIAL UNDER CONTROL

It is possible to get the privileged account threat potential (and, therefore, the threat) under control. The 5 steps in this white paper were designed and presented in a way to ensure you can see the problem for what it is, visualize where your own security is weak, grasp how to address those weak points, and ensure the security controls you put in place over your privileged accounts are under constant watch. By actionably walking through these steps, you begin to create a comprehensive layered approach – and strategy – to stop threats at each of the three layers of an attack.

And, by putting these steps into practice, you're preparing your organization for the continued evolution of privileged access threats. Come what may, and you have the tools to see where you're vulnerable and whether you're under attack, the protocols and process to keep privileged accounts used appropriately, and the ability to know if all these security measures have failed.

The bottom line is that the privileged access threat potential is just that – a potential. Do nothing and the potential is high. Follow these steps, and the *potential* lessens significantly.

ABOUT BEYONDTRUST

BeyondTrust is the worldwide leader in Privileged Access Management (PAM), empowering organizations to secure and manage their entire universe of privileges. Our integrated products and platform offer the industry's most advanced PAM solution, enabling organizations to quickly shrink their attack surface across traditional, cloud and hybrid environments.

The BeyondTrust Universal Privilege Management approach secures and protects privileges across passwords, endpoints, and access, giving organizations the visibility and control they need to reduce risk, achieve compliance, and boost operational performance. Our products enable the right level of privileges for just the time needed, creating a frictionless experience for users that enhances productivity.

With a heritage of innovation and a staunch commitment to customers, BeyondTrust solutions are easy to deploy, manage, and scale as businesses evolve. We are trusted by 20,000 customers, including 78 of the Fortune 100, and a global partner network. Learn more at www.beyondtrust.com.