



ConversationalGeek

ACTIVE DIRECTORY

Security Risk Assessment and Attack Path Management

By Brien M. Posey, Microsoft MVP

While it is sometimes easy to fall into the trap of thinking of Active Directory purely from an administrative standpoint, it also needs to be an area of major emphasis for security professionals. That's because Active Directory has long been a favorite target for attackers. Your adversaries know that if they can take over your Active Directory then they essentially own your enterprise. Unfortunately, many of the security best practices pertaining to Active Directory are completely inadequate for stopping the most common types of attacks.

How Does an Active Directory Attack Work?

Most of the attacks targeting Active Directory begin with a single compromised user account. An attacker may use malware to steal a user's credentials, or they may harvest a user's credentials from one of the many leaked password databases that exist on the Dark Web.

Once an attacker has a set of user credentials that they can use, the next step is to try to gain elevated permissions by gaining access to an administrative account. There are many different ways that an attacker can potentially accomplish this goal, as documented by the MITRE Matrices

(<https://attack.mitre.org/matrices/enterprise/windows/>).

However, if the ultimate goal is to attack the Active Directory, then one of the most common attack methods is to use a Pass-the-Hash Attack.

Once the attacker has a set of user credentials that they can use, they set about attempting to remotely log into the various devices on the network. Upon logging into each machine, the attacker extracts password hashes from the machine's Security Accounts Manager. The Security Accounts Manager contains a hash (essentially a set of cached credentials) for anyone who has ever logged into the machine. The attacker's goal at this point is to find a hash for an account that has administrative permissions. If no such account exists within a machine's Security Account Manager database, then the attacker will move laterally from one machine to another until they eventually discover a machine that contains the credentials that they need. The reason why attackers use this method is because it does not require them to crack an admin password. The stored hash can be used as a password, without the actual password needing to be known.



Why Current Active Directory Defenses Are Not Good Enough

There are countless reasons why mainstream Active Directory defenses are insufficient. These reasons range from the highly dynamic nature of Active Directory to changes in the strategies that attackers use against Active Directory environments.

One of the main reasons why mainstream Active Directory defense techniques are often inadequate is because they fail to address Configuration Debt. This essentially refers to the idea that an organization's Active Directory may have been misconfigured many years ago and the organization simply lives with the problem – either knowingly or unknowingly.

Another issue is the opaque nature of Active Directory permissions. Active Directory complexity can make it difficult to definitively determine who has access to what.



Attackers Are Doing Things Differently

At one time Windows operating systems had a reputation for being extremely insecure. In fact, Microsoft took the extraordinary step of pausing development work on Windows Vista in order to focus on addressing some of the most egregious vulnerabilities that existed in Windows XP.

Being that Windows used to be so insecure, attackers routinely focused their efforts on exploiting well-documented vulnerabilities within the operating system. As time has gone on, however, Microsoft has done a much better job of addressing vulnerabilities and making Windows much more secure. Yes, vulnerabilities exist even in Windows Server 2022 and in Windows 11, but Microsoft typically releases security patches that remove vulnerabilities soon after they are discovered.



There are some attackers who still look for unpatched systems in an effort to exploit vulnerabilities that exist in the absence of security patches. However, attackers know that patch management has been around for long enough that most admins understand the importance of applying security patches in a timely manner. They also know that once a vulnerability is discovered, they have a very short window for being able to figure out how to exploit the vulnerability and then launching attacks against targets. Microsoft's next Patch Tuesday will likely remove the vulnerability, meaning that the work that the attacker has put into figuring out how to exploit the vulnerability might have been for nothing.

The pace with which Microsoft releases security patches has made it increasingly impractical for cyber criminals to develop attacks that are based around exploiting known vulnerabilities. Yes, there will always be those who seek to attack vulnerable systems. By and large, however, attackers have changed tactics and have begun looking for architectural weaknesses that can be exploited rather than focusing on bugs and security vulnerabilities.



From an attacker's standpoint, there are numerous benefits to using this approach. For one thing, because the attacker is not attempting to exploit a bug or a vulnerability, the attacker does not have to break anything. This potentially means that there is less work involved for the attacker and less chance that they will be caught. It also means that the attacker does not have to worry about their hard work being undone by a security patch. After all, system architectures tend to be semi-permanent, even as the underlying operating systems evolve.

So, consider this new philosophy from the standpoint of an attacker who wants to hack an organization's Active Directory. That attacker knows a couple of things. First, they know that at some point along the way, a well-intentioned admin likely tried to secure Active Directory without fully understanding the implications of their actions, and that the permissions that were put into place may actually benefit the attacker.



Second, the attacker knows that even if an organization's domain controllers are running Windows Server 2022, which offers the latest and greatest in terms of security features, attackers know that these new advances are largely meaningless. There are a few different reasons for this.

First, unless an organization is a startup, their Active Directory is probably much older than the domain controllers on which the directory service is running. Active Directory was introduced in Windows 2000, which was released in 1999 – well over 20 years ago. Additionally, a lot of organizations had Windows NT environments in place long before the Active Directory was ever conceived. Windows NT dates back all the way to 1993.

All of this is to say that an organization may have started out with a Windows NT domain, upgraded to a Windows 2000 based Active Directory, and then performed subsequent upgrades as new versions of Windows were released. Even if such an organization is running Windows Server 2022 based domain controllers, their directory structure could be up to 30 years old!



Even though this idea might sound ridiculous, it is important to remember that enterprise class organizations tend to be extremely risk averse. Outages can be extremely expensive, and no administrator wants to be the one to make a major permission change to Active Directory that ends up causing unforeseen problems. Admins often make a calculated judgement in which they weigh the potential security benefits of such a change against the risk of unforeseen breakage. The result of such deliberations is often that admins decide to play it safe by not making the change.

What this means is that an organization's Active Directory structure likely has not kept pace with Microsoft's evolving best practices. In fact, there is a good chance that a misconfiguration that was put in place many years ago may still exist. Many organizations are not even using the various Active Directory security tools that Microsoft makes available.

All of these factors collectively leave Active Directory weak and vulnerable, which makes an attacker's job much easier.

Understanding Technical Debt

Some examples of misconfigurations due to decades of technical debt include:



Multi-level group nesting without any form on ongoing owner attestation - this can lead to large numbers of users with unchecked privileges admins aren't aware of and the users don't need.

Not understanding what Tier 0 actual is, or which systems and objects belong in it.



Kerberos delegation on Tier Zero objects - Active Directory provides several mechanisms to offload Kerberos authentication to other principals. Those mechanisms include two flavors of Kerberos delegation: constrained and unconstrained. Attackers may abuse both flavors of delegation to impersonate other users in the domain and escalate their privileges.

During the delegation process, tickets are created by principals trusted to perform constrained or unconstrained delegation; however, tickets created for other users are only valid if the user is not marked as "sensitive and cannot be delegated" or added to the "Protected Users" group in Active Directory. Users marked in this way or added to that group are effectively "immune" from Kerberos delegation attacks.



Kerberoastable privileges accounts - where the account has values in its serviceprincipalnames attributes and/or does not have a strong password (e.g. at least 32 characters)

Assigning large default groups Generic All privileges - The "Generic All" privilege grants principals the ability to perform nearly any action against the object, including abusable actions such as resetting user passwords and adding new users to security groups. Generic All is also known as Full Control.

Three default groups within Active Directory (*Domain Users, Authenticated Users, Everyone*) should arguably never be granted this privilege, as this is against best practice for least privilege. When an adversary discovers an object where most or all domain authenticated users have this privilege, such objects can provide the attacker with their first step in a critical attack path leading to the compromise of the entire enterprise.



Assigning large default groups in RDP Users group - The Windows operating system grants remote execution privilege to principals that belong to the Remote Desktop (RDP) Users groups. Additionally, Windows grants remote execution privilege to principals that belong to Active Directory security groups that have been added to the aforementioned local groups. Three default groups within Active Directory (*Domain Users, Authenticated Users, Everyone*) should arguably never be members of local groups on any machine as this is against best practice for least privilege. When an adversary discovers a system where most or all domain authenticated users have remote execution and/or admin rights, such systems can provide the attacker with their first step in a critical attack path leading to the compromise of the entire enterprise.

Risk Assessments GONE WRONG

Outdated Active Directory architectures and misconfigurations can be problematic from a security standpoint, but those aren't the only issues. Another major problem is that the general approach that organizations take with regard to risk assessment and Active Directory security also tends to be flawed.

The IT industry has conditioned security professionals to think in terms of lists. For example, there are lists of the security patches that need to be applied to every system. There are lists of group policy settings that need to be put into place. There are lists of security practices that a given organization must adhere to in order to maintain their regulatory compliance. These are just a few examples. The point is that security pros tend to think in terms of lists. There is often a subconscious belief that "if I just do these 25 things, then the organization's IT resources will be secure".

While checklists do have their place, they should be used as an aide, not as a substitute for true security. The problem with relying on lists is that the cyber criminals aren't using them. No self-respecting hacker begins an attack by comparing an organization's defenses against a list of security best practices. That's just not how hackers operate.

Hackers tend to think in terms of graphs, not lists. A hacker's goal is to map an attack path that will get them from a point of entry to their ultimate target (such as an Active Directory takeover). The hacker's graph maps out that attack path for the hacker. All of this is to say that penetration testing, while important, won't fully address Active Directory security risks because penetration testing focuses on individual vulnerabilities rather than on attack paths.

While checklists do have their place, they should be used as an aide,

not as a substitute for true security.



A Better Approach

If an organization is to truly secure its Active Directory environment, it will need to adopt a next generation approach to security that is significantly different from what has been used in the past.

The first thing that IT pros will need to do is to take a frank and honest look at their Active Directory structure. The goal should be to identify Tier 0 assets (as well as who has access to those assets), and then figure out how best to structure the Active Directory environment in an effort to keep those assets secure. Tier 0 assets are those resources that are the most critical to the Active Directory's health and well-being. It includes things like domain controllers, Azure AD Connect servers, DNS servers, domain level or sensitive GPOs, built-in admin groups like Enterprise Admins, service accounts, and servers where any AD backups are housed and any other resource that if compromised, would allow for the takeover of the Active Directory.

As organizations work to protect their Tier 0 resources, they should focus less on security checklists and more on attack path management. In other words, organizations must consider how they can cut off choke points to Tier 0 assets in a way that will block a potential attack path. Of course, in order to do this, organizations must work to identify any potential attack paths that may exist that would ultimately allow an attacker to sign in as a domain user and eventually gain domain admin permissions.

Another important consideration is that many organizations perform detailed risk assessments a couple of times a year. But, when it comes to Active Directory, these types of periodic risk assessments are inadequate. Remember, Active Directory is a highly dynamic environment. New Active Directory objects are created every single day. The simple act of creating a new object or moving an existing object could potentially open up an attack path that did not exist previously.

If an organization is to keep its Active Directory secure and protect against attacks that target the Active Directory structure, then it must adopt a protective strategy that mimics that of cyber criminals. Just as attackers are continuously looking for attack paths, organizations too must be constantly on the lookout for viable attack paths that could be exploited by hackers in an effort to take over their Active Directory. As Active Directory is constantly changing, this requires a real-time monitoring solution that is specifically designed to look for potential attack paths.

Although real-time monitoring is critical to defending Active Directory, monitoring alone is not enough. If it were, then there would be no problem with continuing to use the old detect and recovery strategy that we have all relied on for years. Rather than focusing squarely on monitoring, organizations should seek to implement a comprehensive, end-to-end approach to Active Directory security that encompasses risk assessment, continuous monitoring, mitigation, auditing, and disaster recovery.

As a final thought, there is an old saying that if you can't measure something then you can't secure it. Consider for example, the use of least privileged management (LPM). LPM is an important part of defense in depth, but it can be hard to quantify its impact and it requires a lot of time and resource. Any Active Directory security solution that an organization puts in place should offer quantifiable metrics. It's important to be able to see how the IT department's efforts improve security over time in a quantifiable way.



Quest[®]

Quest creates software solutions that make the benefits of new technology real in an increasingly complex IT landscape. From database and systems management, to Active Directory and Office 365 management, and cyber security resilience, Quest helps customers solve their next IT challenge now. Around the globe, more than 130,000 companies and 95% of the Fortune 500 count on Quest to deliver proactive management and monitoring for the next enterprise initiative, find the next solution for complex Microsoft challenges and stay ahead of the next threat. Quest Software. Where next meets now.

For more information, visit: www.quest.com