How to stop attackers from owning your Active Directory

Received (in revised form): 5th February, 2022

Carolyn Crandall
Chief Security Advocate, Attivo Networks, USA

Carolyn Crandall is the Chief Security Advocate at Attivo Networks, leader in identity detection and response solutions. She is a high-impact technology executive with over 30 years' experience in building new markets and successful enterprise infrastructure companies. Carolyn has a demonstrated track record of taking companies from pre-IPO through to multibillion-dollar sales and has held leadership positions at Cisco, Juniper Networks, Nimble Storage, Riverbed and Seagate.

Attivo Networks, 46601 Fremont Blvd, Fremont, CA 94538, USA
Tel: +1 510-623-1000; E-mail: Carolyn@attivonetworks.com

Tony Cole
Chief Technology Officer, Attivo Networks, USA

Tony Cole has more than 35 years' experience in cyber security and today is the Chief Technology Officer at Attivo Networks, responsible for strategy and vision. Prior to joining Attivo Networks, he served in executive roles at FireEye, McAfee and Symantec and is a retired cyber operator from the US Army. Tony previously served on the NASA Advisory Council and the (ISC)² Board of Directors as Treasurer and Chair of Audit and Risk. Today he serves on the Gula Tech Foundation Grant Advisory Board, helping the foundation give back to the community to drive a more diverse cyber workforce.

Attivo Networks, 46601 Fremont Blvd, Fremont, CA 94538, USA
Tel: +1 510-623-1000; E-mail: Tony@attivonetworks.com

Abstract More than 90 per cent of organisations use Active Directory (AD) as their identity management system, which serves as a master directory and the means to control access to enterprise services. Its central role in governing user identity and authentication means AD is a primary target for threat actors. Compromising AD means attackers can access the most critical systems and assets on the network or gain administrator privileges to take over the domain. Many traditional security solutions will not notice this activity because the user account appears to be operating within the scope of its privileged access rights. The tactics the attackers use can evade traditional detection systems since they are not designed to detect credential theft, privilege escalation and lateral movement. Identity visibility solutions reduce the attack surface by identifying exposed credentials, domain controller vulnerabilities and cloud overprovisioning. Identity detection and response (IDR) solutions add detection of attempts to exploit AD and credential protection from theft and misuse. This paper will discuss how threat actors attack and exploit AD, and what organisations can do to protect their AD environments.

KEYWORDS: Active Directory protection, cyber deception, credential protection, identity detection and response (IDR), identity security, domain controller attacks, ransomware preparedness

INTRODUCTION

In today's digital world, several solutions and assets have become truly indispensable. Active Directory (AD) is one of the most critical of all. Look behind the scenes of any organisation, and the chances are high that AD is there in the background underpinning everything they do. Indeed, around 90 per cent of Fortune 1000 companies rely on it.[1]

AD is essential because it facilitates our fundamental ability to work in the modern office environment. It includes a database and set of related services that enable users to connect with network resources, and it is responsible for critical authentication and authorisation processes across enterprise resources.

Its central role in governing user identity and authentication means AD is a primary target for threat actors. Gaining access to the database will grant a cyber attacker a huge advantage in different malicious activities. Compromising AD enables a threat actor to change access control lists (ACL), group membership, permissions and security policies with impunity. The most direct impact is usually to change the rights of another user account they have taken control of to facilitate lateral movement, privilege escalation and ultimately, the access to steal data or cause disruption to services.

Compromising AD means attackers can access the most critical systems and assets on the network or gain administrator privileges to take over the domain (see Figure 1). Many traditional security solutions will not notice this activity because the user account appears to be operating within the scope of its new privileged access rights. The tactics the attackers use can evade most detection. Accordingly, attacks that exploit AD benefit from long dwell times that enable intruders to deal maximum damage before they are detected.

Armed with this power, attackers can do almost anything, from establishing persistence through backdoor access to striking with precision-targeted ransomware that will begin by encrypting the organisation's most vital assets.

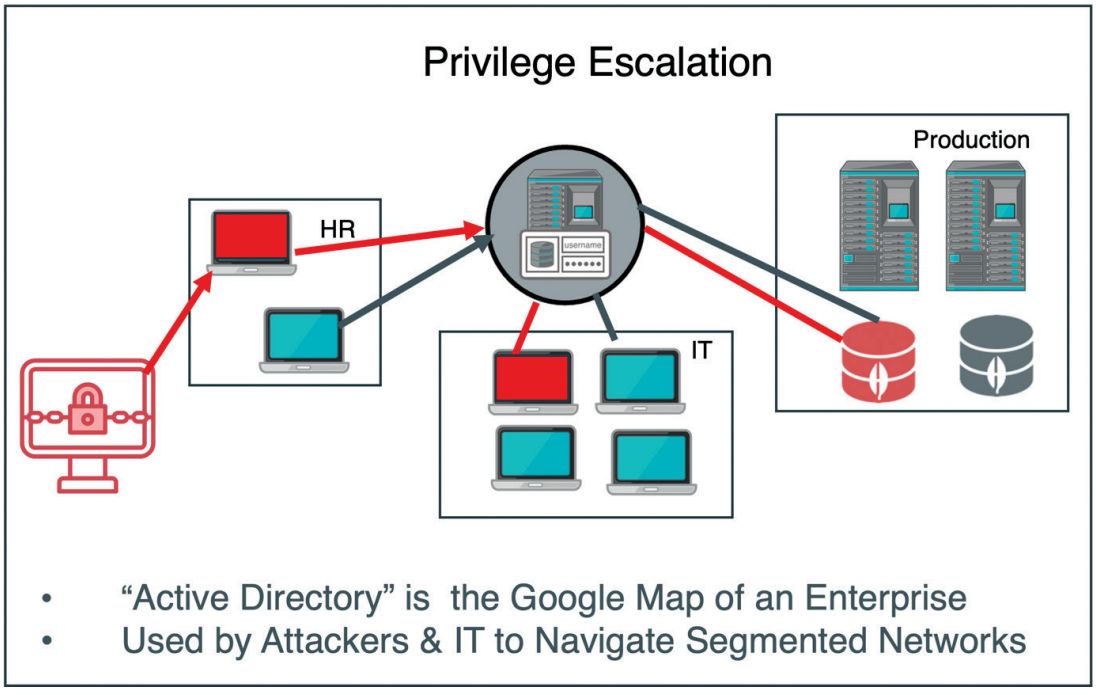

Figure 1: Privilege escalation

This threat means that protecting AD should be a leading priority of any organisation's security strategy; however, one often finds that enterprises have left their AD far too vulnerable to attack.

WHY IS AD OFTEN EXPOSED?

AD is so intrinsic to a functional business that many often treat it more akin to operational plumbing. It is something to be installed and then left well alone as much as possible in case tinkering with it accidentally breaks something.

Performance is measured primarily on service availability, with security being a secondary concern. The size and intricacy of AD mean effective management demands significant expertise and resources, and its complex nature means it is easy to miss vulnerabilities and security gaps. The complexity of the environment will also multiply exponentially when mergers and acquisitions are also factored in.

AD defence has traditionally focused on applying security updates to manage vulnerabilities. Companies will also often employ tiered administration policies, which serve to separate the application into isolated sections. Alongside this, there will usually be an attempt to follow a 'least privileges' approach, with AD only accessible by user accounts that have a genuine need to do so in their job role.

While these strategies are essential, they are no longer enough to keep up with increasingly innovative and aggressive threat actors. Least privileged access should be a standard approach, but on its own, it will only slow attackers down, as they can gradually escalate until they claim an account with the required privilege. Least privilege access can also be difficult to maintain, as AD users change roles across the organisation and gain additional privileges, while retaining existing ones they may no longer need.

Likewise, a good patching cadence, for example, will close most known weaknesses and ensure there is no low-hanging fruit. Still, it can only account for vulnerabilities that have already been discovered and have had patches issued.

This kind of reactive thinking leaves organisations vulnerable to attacks using more sophisticated techniques. Even log analysis and security information and event management (SIEM) correlation activities that identify malicious behaviour tend to be more about reactive detection than proactively searching out system exposures and misconfigurations that create vulnerabilities.

The problem appears to be increasing in severity, with the 'Verizon 2021 Data Breach Investigations Report'[2] stating that 61 per cent of all breaches now involve credential data, including stolen credentials, credential stuffing, brute force attacks, credential leaks and more. This validates that all too many organisations lack sufficient visibility into exposed credentials, particularly as the attack surface has expanded amid widespread remote work — further highlighting the need for effective visibility, prevention and detection solutions.

Almost every ransomware attack has leveraged AD to gain the control and access required to distribute ransomware and encrypt systems. The EMA 2021 survey on AD protection[3] states that 50 per cent of all respondents had experienced an attack on AD in the last one to two years, with 40 per cent of the attacks being successful. Again, this reinforces the need for organisations to look beyond traditional defences and seek ways to improve identity exposure visibility and identity detection and response (IDR) controls.

HOW DO THREAT ACTORS ATTACK AND EXPLOIT AD?

As with most other forms of cyberattack, strikes targeting AD usually begin with an initial endpoint compromise, likely carried out via phishing. Once the adversary has

claimed an initial set of stolen credentials and accessed the user's device, they will begin hunting down further credential sets that will enable them to move laterally and set up back doors that will help them evade detection.

AD plays a critical role here, as attackers can query it to discover who the domain admins are, pointing them in the right direction. They use multiple attack techniques to gain access to AD or exploit its capabilities, including Windows Security Identifier (SID) history injection, 'golden ticket' attacks and Kerberoasting.

Once adversaries obtain access to a domain controller account, they will have practically limitless power within the system (see Figure 2). This is often referred to as a game-over situation and the targeted company is destined for a seriously damaging and costly breach. The clean-up after this form of attack can often outweigh even the initial ransomware demands, given the risk that the attacker or other cyber foes could return.

The last few years have seen a decline of 'smash and grab' attacks in favour of slower, more advanced persistent threat (APT)-style targeted strikes. Most criminals would previously have rushed to access and exfiltrate as much data as they could reach after breaking into the network before they were inevitably discovered and stopped. While these attacks do still occur, we are seeing an increase in more deliberate intrusions, with the adversary taking the time to slowly achieve lateral movement and privilege escalation to hunt out the most valuable assets. AD's central role in identity and access management makes it the first port of call for these low-and-slow attacks.

Criminals are also increasingly combining these more sophisticated data exfiltration attacks with ransomware, dubbed 'Ransomware 2.0'. Much like smash-and-grab-style data breaches, ransomware was largely a blunt instrument in previous years. The malware frequently infiltrated through phishing e-mails and websites and, upon activation, would immediately begin encrypting anything within reach. Companies that lacked effective network segmentation or threat detection capabilities would likely still see most of their files encrypted, but more prepared organisations had a good chance of stopping the outbreak before it reached anything important.

With Ransomware 2.0, once the intruders find and steal the most valuable assets, they unleash their ransomware right in the heart of the network, ensuring that the outbreak will cause maximum damage even if the

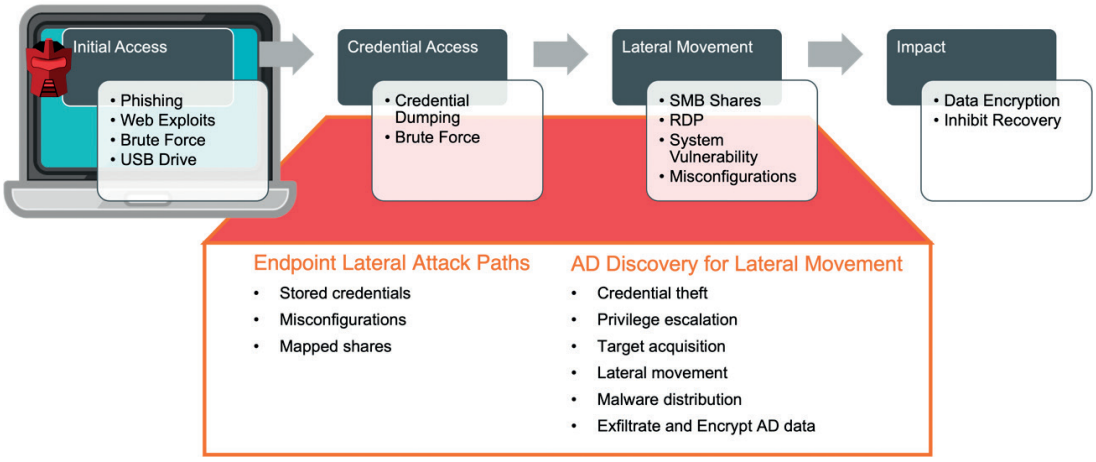

Figure 2: Disrupting the ransomware attack chain

victim can identify and contain it reasonably quickly. Indeed, AD itself is likely to be a key target of Ransomware 2.0 attacks, as locking the system down will deal a crippling blow to the organisation.

Attackers know that their victims will more likely pay the ransom to resolve targeted Ransomware 2.0 attacks that compromise AD and cause maximum disruption. Accordingly, the average ransomware payout has increased by over 171 per cent in the last year alone, with the average cost[4] now exceeding US$300,000. On the higher end of the spectrum, organisations such as JBS USA and Colonial Pipeline paid millions of dollars in ransom to the hackers that launched their attacks. In the case of Colonial Pipeline, the incident all started with the compromise of one credential.

Furthermore, because the attackers also exfiltrate valuable data before encrypting systems, they can still secure a sizeable payday even if the victim refuses to give in to ransom demands. Some groups will also use stolen data to squeeze their victims further by blackmailing customers and threatening to release sensitive information publicly unless they meet additional payment demands.

Ransomware has become a lucrative business for organised crime, given the size of the payouts that can be secured. For example, the Conti ransomware gang has been linked to over 40 attacks in 2021[5] and most recently was linked to an attack on Delta Electronics, a Taiwanese electronics manufacturing company and a major supplier of power components to companies including Apple and Tesla. This is a powerful example of attackers disrupting the supply chain and not limiting their attacks to only well-known brand names. Delta is primarily known for its UPS solutions.

HOW CAN AD BE PROTECTED?

Against these threats, organisations must treat AD as a high-priority attack target because the critical infrastructure keeps their operations running; however, the scope and complexity of the system make it intrinsically insecure and difficult to protect. Because of this, attackers can exploit many attack paths and techniques to access AD and manipulate it without raising the alarm.

As such, AD security strategies need to follow a defence-in-depth approach that includes multiple layers. The focus should start by eliminating low-hanging fruit that will grant easy access and by enabling real-time automated monitoring for exposures from endpoints and within AD. AD cloud deployments must also be factored in, as attackers will jump from on-premises to the cloud and vice versa.

In addition, organisations should focus on security solutions that will allow them to detect credential theft, privilege misuse and attacks on AD. Security teams need to respond to these attacks in real time to shut them down before they cause damage.

Early detection and response capabilities can combine with the ability to misdirect intruders, hiding AD and other critical assets and leading them astray. Deception and concealment technologies can provide invaluable capabilities to detect, misdirect, misinform, deny and derail such attacks.

Organisations can create realistic decoy assets, mimicking real files and systems likely to be targeted in an attack. Ideally, these fakes will not only appear genuine to automated scanning tools attackers use, but will also bear a degree of manual interaction and scrutiny.

The longer an adversary is deceived into wasting their time exploring the decoy network, the greater the chance organisations will discover them and shut them down before they can do any damage. Organisations can configure lures to trigger an alert as soon as attackers interact with them, giving the security team a powerful head start. Alongside shutting down the attack, the security team will have the

opportunity to study the adversary and learn about their tools, tactics, likely targets and motivation. They can then use these insights to harden the organisation's defences against further attacks.

Deceptive tactics work best when combined with multiple layers of defence. The following are some of the most important security priorities, mixing deception with effective management and monitoring.

FIND AND SECURE ALL PRIVILEGED ACCOUNTS

Most cyberattacks begin by exploiting user credentials to evade perimeter-based controls, either by directly stealing them via phishing or simply purchasing them from other criminals on the dark web.

If the organisation has done an excellent job of applying zero trust architectures or a least-privileges approach, the attacker will need to compromise a specific set of accounts to access AD. Enterprises should be making this task as difficult as possible by keeping privileged accounts well secured, but this is not always the case.

Many users will store their credentials on their endpoint devices or in cloud environments for convenience, particularly if they have additional credentials for accounts with privileged access to different systems. Once an attacker compromises an account, they can quickly sniff out any exposed sets of credentials and use them to escalate their privilege. They can also find and exploit credentials that the users have previously sent via e-mail or shared on channels like Teams, along with any credentials stored in virtual private network (VPN) clients, Windows registries, file transfer protocol (FTP) clients or other applications on the compromised endpoint.

To mitigate this risk, enterprises should take action to identify all privileged account exposures, including credential sets saved to endpoint devices, shared folders, or sent via e-mail and other channels. Organisations may also use privileged access management (PAM) for on-premises AD and privileged identity management (PIM) solutions for Azure AD, which provide the means to control, monitor and access privileged accounts. Protection of privileged access is a good best practice, but still leaves opportunity for attackers to leverage machine-to-machine communications. Bad password practices, policy drift and the use of local admin credentials on endpoints or privileged accounts can still leave avenues for attackers to gain privileged access. These exposures are best addressed with credential identity exposure visibility and cloud infrastructure entitlement management (CIEM) tools.

STOP ATTACKERS EXPLOITING DOMAIN SHARES TO HARVEST CREDENTIALS

Alongside credential sets saved in usual formats such as Word documents, spreadsheets and e-mails, adversaries can also access plaintext or reversible passwords in scripts or group policy files. These get frequently saved in domain shares like Netlogon or Sysvol, which are available throughout the domain.

Security teams can use automated tools to identify these password sets rapidly and secure them properly before threat actors discover and exploit them.

Organisations can also protect these assets well by using a deceptive strategy. For example, creating realistic decoys of Sysvol group policy objects in the production AD will misdirect attackers away from the genuine asset and cause them to waste their time trying to access a useless fake. High-quality decoys that can withstand a degree of manual inspection by an attacker or automated scans will maximise the misdirection and greatly increase the chances of the security team stopping the intruders before they can access AD.

DISCOVER AND RESOLVE ALL PRIVILEGED ACCOUNTS WITH DELEGATION

Adversaries can skip directly obtaining privileged credentials by hunting for accounts with delegation enabled. This setting allows AD administrators to delegate capabilities and access privileges to other accounts to execute certain administrative tasks without logging in with separate accounts. Organisations often use this approach to manage AD, as it enables other users to have elevated access privileges without adding them to privileged groups like domain admins and account operators.

While convenient, it can also create a significant security risk, as attackers can readily exploit an account with delegated abilities if they can compromise it. Attackers abuse unrestrained delegation for techniques such as Kerberoasting and silver ticket attacks, explained below.

Enterprises, therefore, need to detect and report on any privileged accounts that have delegation enabled. Being armed with a thorough list of all privileged users and delegated admins will help security teams to get a full view of vulnerabilities that attackers may exploit. Security heads should ensure this process includes service accounts, which often get overlooked, as automated systems rather than humans largely use them.

Account delegation can be extremely useful in managing AD and other key system assets, so organisations using it need to ensure they also have additional layers of security to prevent attackers from discovering and exploiting delegated accounts.

PREVENT THE ENUMERATION OF PRIVILEGED, DELEGATED ADMIN SERVICE AND NETWORK SESSIONS

After an attacker has established a foothold in the system and has managed to gain access to AD, they will quickly exploit it to conduct reconnaissance on valuable assets and how to reach them. They will gather information on various accounts and their capabilities through enumeration, such as acquiring a list of all domain admins.

Enumeration activity can be completely legitimate, and attackers can easily disguise their attack activity as normal business use. Most traditional security solutions will not detect anything wrong if attackers use a captured account with the appropriate access or there are no effective privilege controls in place around AD in the first place.

Suppose security teams can quickly identify unauthorised enumeration of privileged, delegated and service accounts with admin powers. In that case, they have an increased chance of detecting and preventing the attackers from accessing and exploiting these accounts. Discovering an intruder early into the attack cycle increases the likelihood of shutting the attack down before it can cause an actual security breach.

Deceptive tactics can be useful in this scenario, as fake domain accounts and credentials planted on endpoints can redirect attackers and slow them down while simultaneously hiding the actual production credentials and preventing unauthorised access to local credential stores. It can also configure decoys as destinations for the lures to alert security teams as soon as an attacker interacts with them, creating an opportunity to observe the intruder and gain insight into their tactics and goals in addition to preventing damage to production assets.

IDENTIFY AND PROTECT ACCOUNTS WITH PRIVILEGED SID

Attackers may also move laterally within the AD environment and escalate their privileges using the SID injection technique. Windows environments use SIDs to provide a unique value to identify a user or group account and to label access tokens and security descriptors. The SID history Active Directory attribute allows accounts to hold additional SIDs that facilitate interoperability and account migration between domains.

A SID injection attack sees the adversary insert SID values into an account to gain elevated access, such as impersonating a member of the domain admin group.

To mitigate the risk of this technique, organisations should use PowerShell to identify any accounts with well-known privileged SID values proactively. After completing account migration, they should also clean up SID history attributes to reduce their exposure to threat actors. Security teams can also monitor management events to alert them to any attempts to change SID history.

IDENTIFY AND STOP TICKET ATTACKS

Pass-the-ticket (PTT) attacks are among the most effective and dangerous techniques threat actors use to move laterally and escalate privileges. The attack involves the adversary extracting a Kerberos ticket granting ticket (TGT) from the local security authority subsystem service (LSASS). They then use the TGT to request Kerberos ticket granting service (TGS) tickets on another system, granting network access. This technique is particularly problematic because Kerberos's stateless design strategy makes it fairly easy for attackers to forge tickets.

The 'golden ticket' is a particularly infamous variant. Like its Willy Wonka namesake, this technique makes dreams come true — but unfortunately, only the dreams of cybercriminals. With this approach, the attacker gains control of the hidden account responsible for encrypting all other authentication tokens, allowing them to log into any account on the system.

Another variant is the 'silver ticket' approach. Despite the name suggesting it is a runner-up, this attack is even more dangerous and difficult to detect than its golden counterpart since all attack activity is local to the endpoint, with no communication with the domain controller to indicate that an attack is underway.

Dealing with ticket-based attacks requires the ability to detect any vulnerable Kerberos TGT and computer service accounts that attackers could exploit. Organisations should resolve any misconfigurations and consider additional layers of security. Deceptive decoys can again be useful here as they will misdirect attackers and increase the chances of detecting them before they can forge tickets at the endpoint.

DEFEND AGAINST KERBEROASTING, DCSYNC AND DCSHADOW

Adversaries will also seek to exploit Kerberos in an attack type known as Kerberoasting. Here, the attacker extracts service account credential hashes from AD and then cracks them offline to gain privileged access.

Kerberoasting is typically a late-stage attack that occurs when an adversary has already achieved permanence. Once the attacker has gained domain controller-level credentials, they can also execute other attacks such as DCSync, where the attacker impersonates an AD domain controller to obtain credentials from other domain controllers. Similarly, DCShadow sees the attacker use privileged credentials to register a new domain controller to push domain changes.

Effective service account management and strong encryption will play an important role here, as attackers will be looking for poorly secured or overlooked service accounts they can easily hijack. Continuous monitoring and assessment of AD will also help to detect ticket-based attacks.

PREVENT EXPLOITATION OF ADMINSDHOLDER ACL

AD Domain Services (AD DS) secures privileged users and groups through the security descriptor propagation (SDProp) process, periodically propagating changes to built-in AD groups. Alongside this, the AdminSDHolder object uses a unique access control list (ACL) to control the permissions

of security principals that are members of built-in privileged AD groups.

Attackers can exploit the AdminSDHolder ACL by adding accounts that grant them the same privileged access as other protected accounts, thereby gaining lateral movement and access to restricted systems. The SDProp process then propagates this to all other child objects, thereby providing persistence to the change. If a domain admin changes the permissions on a protected group or user, SDProp will change the security permissions to match the AdminSDHolder object.

Enterprises can mitigate the risk of this attack type by implementing solutions that will detect unusual accounts within the AdminSDHolder ACL and alert the security team.

DETECT CHANGES TO DOMAIN POLICIES

Finally, attackers can also exploit AD's capabilities around creating group policies to manage operational configurations. Administrators can configure policies to several key tasks, including setting organisation-defined security requirements at each level, installing software and setting file and registry permissions.

Gaining access to AD means an adversary can change policies to achieve domain persistence, setting them up for several attack types.

Organisations should ensure they can monitor for any changes to default group policies that indicate an attacker's presence.

TAKING A PROACTIVE APPROACH TO KEEP AD SAFE FROM ATTACK

Keeping AD secure is a serious challenge, but not an insurmountable one. Organisations should first prioritise getting the basics right and identifying and resolving commonly overlooked vulnerabilities such as privileged account exposure. Locking down these accounts will deprive threat actors of their easiest attack routes and increase their chances of making a mistake and getting caught out as they seek access to AD.

From here, organisations can move on to IDR solutions, which focus on detecting live attacks targeting AD objects and preventing attackers from making AD changes that grant them control. Innovations in technology will also enhance zero trust security controls by proactively hiding and denying access to AD objects and efficiently redirecting attackers away from their targets. By protecting their AD environments, organisations can now effectively detect lateral movement and the exploiting of AD so that they can shut the attacker down before they can truly strike.

References

1. Krishnamoorthi, S. and Carleton, J. (March 2020), 'Active Directory Holds the Keys to your Kingdom, but is it Secure?', Frost & Sullivan, available at https://www.frost.com/frost-perspectives/active-directory-holds-the-keys-to-your-kingdom-but-is-it-secure/#:~:text=Microsoft%20Active%20Directory%20%28AD%29%20is%20the%20dominant%20mode,primary%20method%20to%20provide%20seamless%20authentication%20and%20authorization (accessed 5th February, 2022).
2. Verizon, '2021 Data Breach Investigation Report', available at https://www.verizon.com/business/resources/reports/dbir/?cmp=knc:bin:ac:ent:security:8003162844&utm_term=data%20breach%20report&utm_medium=cpc&utm_source=bing&utm_campaign=BNG_NB_Security_Phrase&utm_content=Enterprise&msclkid=c728fa1ed553114d54649bea88b9482c&gclid=c728fa1ed553114d54649bea88b9482c&gclsrc=3p.ds (accessed 5th February, 2022).
3. Musich, P. (September 2021), 'EMA Research Report: The Rise of Active Directory Exploits: Is it Time to Sound the Alarm?', Attivo Networks, available at https://go.attivonetworks.com/WC-EMAReportAD_LP-Registration.html?utm_source=EMA+Research&utm_medium=report&utm_campaign=Active+Directory (accessed 5th February, 2022).
4. Wadhwani, S. (March 2021), 'Average Ransomware Payout Touched $312K in 2020, Up From $115K in 2019', Toolbox, available at https://www.toolbox.com/it-security/threat-reports/news/average-ransomware-payout-touched-312k-in-2020-up-from-115k-in-2019/ (accessed 5th February, 2022).
5. HHS Cybersecurity Program (February 2022), 'Lessons Learned from the HSE Cyber Attack', U.S. Department of Health & Human Services, available at https://www.hhs.gov/sites/default/files/lessons-learned-hse-attack.pdf (accessed 5th February, 2022).