Minimizing the Active Directory and Entra ID Attack Surface

Summary

For security practitioners, it's no surprise that the identity management landscape has for decades been dominated by Active Directory. Entra ID, formerly Azure Active Directory, is a popular choice for hybrid environments. Active Directory poses a polarity that organizations' security teams must manage:

1. **Features and capabilities that have made Active Directory the overwhelmingly preferred choice of global organizations**
2. **Vulnerabilities that come from design as well as from human mistakes**.

Vulnerabilities, if unknown or not addressed, open the door to cyberattackers wishing to steal from you or do you harm. It is your responsibility to stop them.

Over 90% of exploits target Active Directory

These dangers lurk within the identity management infrastructure – it is complex and broad, involving authentication and access management for all users, servers, workstations and network devices in your ecosystem. That is easily thousands in many companies. Even with comprehensive monitoring and robust auditing, attackers can still find ways in. Once in, assume they look to spread their attacks far and wide within your organization.

Specific strategies and actionable tactics can minimize the attack vectors that attackers exploit. Taken together, all those vectors make up the "attack surface".

If you are responsible for securing an identity and access management environment, adopting measures to reduce the Active Directory and Entra ID attack surface will significantly benefit you and your organization.

Overview of Active Directory and Entra ID Attack Surface Significance

Identity Management Security Posture – Is Yours as Strong as You Think?

Consider the details of BadSuccessor.

In May 2025, an Active Directory vulnerability and threat, each named “BadSuccessor” by Akamai researchers [https://www.akamai.com/blog/security-research/abusing-dmsa-forprivilege-escalation-in-active-directory] was discovered and quickly in the press. This newly found vulnerability aWecting Windows Server 2025 is caused by a new feature, the delegated Managed Service Account (dMSA), that was introduced by Microsoft as an intended security improvement. The dMSA was meant to allow migration from existing legacy service accounts, but a serious flaw in the implementation can invite privilege escalation attacks.

This vulnerability is:

- Enabled by default in Windows Server 2025
- Trivial to implement, so is an easy target, even for low-skilled attackers
- Widespread and requires commonly held permissions. Permissions required by

BadSuccessor are permissions that 91% of environments (examined by Akamai) have granted to users outside domain admins: ‘CreateChild’ on any Organizational Unit.

BadSuccessor can initiate identity threats in most Active Directory environments that are diWicult to detect. It is easily exploited by linking a dMSA to a predecessor account. The result of that link is that full access to predecessor resources is gained to everything from password hashes to directory replication. If an organization lacks dMSA event monitoring (many do), exploitation of BadSuccessor may go unnoticed.

The dMSA feature has created a wider attack surface; there is currently no patch for this vulnerability, so proactive defensive steps to mitigate have enormous importance.

And this is just one example. Please read on!

Active Directory’s history makes it a prime target

The central place Active Directory has had in identity and access management in most organizations has made it a prime target for cyberattacks. Entra ID is under constant potential threat as well. The challenges to keeping these environments safe from identity thefts, deliberate harm, and ransoms are ever-changing and are not easy to combat, as evidenced by the statistics and stories in the news.

Here are some sobering statistics from the Microsoft Digital Defense Report 2024:

- 88% of customers impacted by incidents had “insecure Active Directory configuration”.
- Median time for an attacker to begin moving laterally after compromise is 1 hour, 42 minutes.
- 68% of organizations impacted by cyber incidents had no eWective vulnerability and patch management process.
- Over 90% of exploits target Active Directory

It’s likely that Active Directory and, more recently, Entra ID have been integral to your organization’s technical Infrastructure,

applications, and data protection for years or decades. Sprawl and excessive privileges happen, and they cause problems that are not easily detected. Vulnerabilities are born of the design of Active Directory and Entra ID, human action (or inaction), and configuration errors. Attackers know where vulnerabilities likely exist in Active Directory and Entra ID environments and take advantage of them. Many organizations have years of Active Directory configurations. Misconfigurations can accumulate over time because of competing priorities and inconsistent enforcement of policies and best practices. For example, avoiding excessive permissions for accounts is a commonly recognized best practice for a strong security posture. But lack of knowledge or visibility or time to devote to the investigation mean that unnecessary elevated permissions on accounts can leave open vulnerabilities to threats against your data and applications.

The results of Active Directory and Entra ID vulnerabilities and exploits by malicious actors include:

- Data theft
- Ransoms
- Business disruptions
- Reputational damage
- High costs of recovery from a breach

This paper will discuss all of these and oWer suggestions on how to minimize the risks of them happening to your organization. First, we examine some real-life examples for clarity of the potential breadth and impact of a cyberattack.

Newsworthy breaches involving Active Directory and Entra ID

Let's look at examples of the disastrous eWects of recent infiltrations and resulting exploits by malicious actors on real businesses.

Significant data breaches involving Active Directory and Entra ID infrastructure over the past two years are highlighted here because of their **resulting large scale of the compromise or high financial losses.**

1. Change Healthcare (February 2024) – The Largest Healthcare Breach in History

One of the largest healthcare payment processing companies in the United States suWered a ransomware attack that exploited their Active Directory infrastructure [1,3]. The attackers, a ransomware group called ALPHV/BlackCat, gained access through compromised credentials and the absence of multi-layer authentication (MFA) on critical systems[4].

The 190 million Americans aWected make it the largest healthcare data breach in U.S. history[3]. The stolen data included:

- Medical records, diagnoses, and treatment information.
- Social Security numbers and driver's license numbers
- Financial and banking records
- Health insurance information and member numbers

The total cost of the breach reached $3.1 billion in 2024, making it one of the most expensive data breaches ever recorded[3][4]. The total cost includes:

- **$1.521 billion in direct response costs**
- **$2.457 billion in total cyberattack impacts**
- **$22 million ransom payment to attackers**
- **Microsoft Midnight Blizzard Attack (2024)**

Microsoft faced an ongoing attack in 2024 from a Russia-based group called Midnight Blizzard, which targeted Microsoft's corporate email systems and extended to customer environments. The attackers successfully exploited vulnerabilities in Microsoft's identity infrastructure to gain persistent access. The attack demonstrated sophisticated techniques that targeted Active Directory, and perhaps other internal systems[1].

The exact number of compromised customers has not been disclosed.

Nor has Microsoft disclosed the financial impact of the Midnight Blizzard attack, but the company faces ongoing costs into 2025 due to incident response, system hardening eWorts, and potential regulatory fines[1].

2. Microsoft Midnight Blizzard Attack (2024)

Microsoft faced an ongoing attack in 2024 from a Russia-based group called Midnight Blizzard, which targeted Microsoft's corporate email systems and extended to customer environments. The attackers successfully exploited vulnerabilities in Microsoft's identity infrastructure to gain persistent access. The attack demonstrated sophisticated techniques that targeted Active Directory, and perhaps other internal systems[1].

The exact number of compromised customers has not been disclosed.

Nor has Microsoft disclosed the financial impact of the Midnight Blizzard attack, but the company faces ongoing costs into 2025 due to incident response, system hardening eWorts, and potential regulatory fines[1].

3. MGM Resorts and Caesars Entertainment Attacks (September 2023)

Both companies were victims of cyberattacks by the group Scattered Spider. Social engineering tactics were exploited to compromise Active Directory credentials[5][6]. The attackers used LinkedIn to identify employees and then called IT help desks to gain network access.

The Caesars Entertainment loyalty program database containing driver's license and Social Security numbers of millions of Americans was compromised[5]. $15 million in ransom was paid to the attackers.

MGM Resorts suWered massive operational disruptions across casino operations[6] but refused to pay ransom. They suWered significant losses from system downtime[6].

4. NHS UK Synnovis Ransomware Attack (June 2024)

The cybercrime group Qilin targeted this blood test management organization, stealing and publishing nearly 400GB of private patient data on the dark web. A ransom of $50 million was demanded by Qilin, but Synnovis refused to pay it[7].

Seven NHS hospitals saw significant disruptions in operations, with more than 3,000 appointments thrown into chaos[7]. Names, dates of birth, and blood test details were stolen.

5. Ascension Healthcare Ransomware Attack (May 2024)

A large nonprofit healthcare system experienced a ransomware attack, the first of two major data breaches in 2024. They highlight the operational and privacy impacts that data breaches can cause.

A Black Basta ransomware attack disrupted clinical operations across all 142 Ascension hospitals. Systems rendered inaccessible included:

- Electronic health records
- Prescription writing
- Scheduling
- Phones

The disruptions meant that ambulances were diverted from hospitals elective procedures were postponed and staW had to revert to paper-based processes [8].

Credential theft (phishing) tactic: The breach began when an employee downloaded a malicious file, thinking it was legitimate. The attackers got a foothold and gained lateral movement within Ascension's network.

The attackers accessed and exfiltrated files from 7 out of 25,000 servers, and it appeared to Ascension that full patient records had not been compromised. But some files did contain protected health information (PHI) and personally identifiable information (PII) [8]. The breach ultimately aWected about 5.6 million patients.

This attack highlighted the spread of impact often associated with data breaches. In this case, as the attackers compromised more systems, patient care was aWected, and uncertainty of the length of delay for care caused significant ongoing operational challenges.

Do you have defenses for attacks like these? Can you mitigate the risks?

How can BadSuccessor threats and exploits be mitigated

To illustrate how mitigation is possible, even with BadSuccessor with its potential of spreading threats and exploits within an organization, here are some specific actions to take:

- Auditing Organization Unit permissions. Unnecessary 'CreateChild' rights can then be revoked.
- Monitoring dMSA object changes. Track changes to 'msDSManagedAccountPrecededByLink', for example.
- Restricting dMSA usage. If they are not required, disable dMSAs.

The attack surface must be reduced to eliminate as many risks as possible. Stopping these attacks in the first place is the desired strategy. And there are some best practices to apply to this eWort.

What Active Directory and Entra ID Attacks Mean to Your Business

The high cost of data breaches

Data breaches are expensive.

When organizations are hit with high costs because of those breaches, they often pass along the costs to their own customers by hiking prices. It's a risky move in markets already facing pricing pressure. Most organizations said in 2024 that they planned to increase prices of goods and services following a breach.

Business disruptions

According to the IBM Cost of a Data Breach 2024, 70% of organizations experienced a significant of very significant disruption to business because of a breach.

Reputational damage

Can you put a price tag on it? A data breach brings costs beyond the eWort to fix vulnerabilities or the payment of a ransom. Business can be aWected greatly by a data breach — customer trust, once shattered, is very diWicult to restore.

Financial costs

Unplanned (sometimes very large) costs are major concerns to companies, and to the security practitioners tasked with protecting sensitive resources, are costs of data breaches.

Financial costs include:

- Business impacts, financial losses
- Direct costs of remedies
- Ransom payments
- Regulatory penalties

According to the IBM Cost of a Data Breach 2024:

- The global average cost of a data breach: $4.88 million
- Average cost of a malicious insider attack: $4.99 million
- Days to identify and contain breaches involving stolen credentials: 292 (the longest of any attack vector)

Researchers found data breaches with a lifecycle exceeding 200 days had the highest average cost, at $5.46 million. The longer it takes to identify to detect and recover from a breach increases the cost of that breach. Finding the cause of the breach quickly so recovery can start often means dealing with diWicult challenges, including security staWing shortages, and building visibility to vulnerabilities that exist in the environment.

Security stacking shortage

Adding to the pressure on security teams in many organizations is the growing staWing shortage and security skills gap. As the threat landscape broadens, the security skills gap is growing. Security team members require expertise in recognizing and dealing with new vulnerabilities and attack techniques. With staWing shortages, time for learning and adjusting is often just insuWicient. The IBM Cost of a Data Breach Report 2024 reports that the growing security staWing shortage increased by 26.2% in 2024 over 2023. More than half of breached organizations are facing high levels of staWing shortage in the security domain. The same IBM report indicates that the average added cost of a data breach due to the security staWing shortage is **$1.76 million**.

Lack of visibility into vulnerabilities — firefighting mode alone means the attackers are winning

Proactive visibility into vulnerabilities before incidents occur is a fixture on most security practitioners' wish lists. But so often, the lack of visibility has contributed greatly to the accumulation of dozens or hundreds of vulnerabilities that threat actors can easily exploit. Tools are critical in providing the visibility you need, and to fully implement the actionable steps this paper describes.

The Active Directory and Entra ID Attack Surface: Risk Factors

Active Directory (and Entra ID in hybrid or cloud-only environments) is integral to most enterprise networks, controlling access to critical resources. And these systems act as the central identity store. So Active Directory or Entra ID compromise isn't just a problem – it's highly problematic, sometimes catastrophic, for the business.

Increasing sophistication of attacks

Attacks targeting Active Directory, e.g., Kerberoasting, Golden Ticket, Pass-the- Hash, DC Sync, are becoming ever more sophisticated and diWicult to detect. The "Attack Surface" of Active Directory – the comprehensive set of vectors of attack initiation - includes misconfigurations, unpatched vulnerabilities, excessive privileges, weak authentication, legacy systems, unmonitored changes, etc. Reducing the attack surface is essential to combat threats as attackers' methods evolve.

Traditional security measures are dieicult and fall short

Privileged Identity Management (PIM) and Privileged Access Management (PAM) solutions don't completely protect from exploitation of privileged accounts, service accounts, and legacy credentials. Misconfigurations or gaps in coverage can allow attackers to escalate privileges or activate admin roles that have been dormant but never removed. All these activities may be undetected.

Antivirus and endpoint protection focus on malware, so they don't protect against identity-centric attacks. Credential theft can still happen undetected, as can privilege escalation or lateral movement once an attacker has breached security measures.

Security scans that are done on some cadences are helpful but can miss active threats. Ever-evolving threats (e.g., Kerberoasting and DCSync) or new types of attacks, can only be detected using continuous monitoring. Monitoring was mentioned in our discussion of the BadSuccessor vulnerability and its mitigation.

Sifting through immense event logs to locate relevant information about what is happening and whether there is suspicious behavior is not productive and certainly isn't a proactive approach. Real-time monitoring and alerting of the entire hybrid identity management ecosystem is needed to detect and respond quickly to identity-based attacks.

Configurations involving default settings that allow all users to add workstations to a domain, or that allow too many users in privileged groups, create vulnerabilities if traditional controls don't detect them before an attack.

People internal to your organization can take actions detrimental to a secured identity system. Phishing and social engineering can target Active Directory credentials. And insider threats, whether negligent or malicious, can mean unapproved changes to configurations.

Hybrid environments are becoming more prevalent, and with them come new challenges. Attackers exploit the lack of unified security across on-premises Active Directory and cloud-based Entra ID. Savvy attackers can bypass defenses because of gaps between the systems.

Attack surface hardening: Thwart these common attack vectors in Active Directory and Entra ID

Some attack vectors that commonly oWer cybercriminals ways into Active Directory and Entra ID – and into the organization's digital assets – can be anticipated. This discussion will provide context and strategies to reduce the attack surface, and actionable steps to take to proactively or correctively defend against attackers.

- **Credential theft**

A wide range of vulnerabilities can leave you open to credential theft attacks like Pass-the-Hash, Kerberoasting, Password Spraying, DCSync, or logons to unsecured computers with privileged accounts, to name a few. Strong password policies, least privilege access, security administrative practices, and continuous monitoring are all defenses that lower the risks of these attacks.

Opening the door to these types of attacks is commonly accomplished using social engineering – phishing – to trick a person into somehow revealing their login credentials. Mitigation starts with something you control: employee training. To minimize the risks of someone being fooled into clicking a malicious link or navigating to an impersonated web page where they might be fooled into providing their corporate credentials.

- **Privilege escalation**

A core risk in Active Directory and Entra ID and probably the fastest way for attackers to take over the identity security system. Attackers exploit known weaknesses in the software or weak configurations like unconstrained delegation and misconfigured ACLs to gain higher privileges like Domain Admin. Techniques they might use include joining administrative groups, extracting Kerberos tickets and highjacking credential to bypass authentication, manipulating Group Policy Objects to abuse delegation or to plant backdoors, among many others.

Entra ID account takeover risks: Identity threats through compromised overprovisioned accounts have become prevalent because of vulnerabilities like Synced Service Accounts that allow attackers to pivot to Entra ID, and misconfigured hybrid identities that lack MFA, for example. Some Entra ID connection exploits (synchjacking, hard matching) are still possible, despite Microsoft changes to use application registrations to reduce risk.

Given the breadth of the vulnerabilities and the risks to attacks of privilege escalation, a broad defense is needed. The strategies and actions listed in the next section will contribute to attack surface hardening.

- **Lateral movement**

These are post-breach tactics that allow attackers to make things worse for the victim by targeting high-value assets like domain controllers. RBCD and GPO abuse are dangers in Active Directory. DNS manipulation is a tactic that bypasses network controls and allows the attacker to move undetected.

Again, continuous monitoring and real-time alerts are keys to minimizing risks, for example, watching for unusual login patterns, unexpected group membership changes, or DNS record modifications. Alert notifications are critical for events like administrative account changes, suspicious delegation, or replication requests.

- **Cloud pivoting**

Pivoting, or transitioning, from on-premises Active Directory to cloud-based Entra ID can expand an attacker's access greatly for identity threats. By exploiting synchronization mechanisms and trust relationships between hybrid environments, attackers can mount potentially broad Entra ID attacks, including:

- Pass-the-PRT: Steal tokens from Active Directory to authenticate to Entra ID resources without MFA.
- Entra Connect Sync Exploitation: Extract credentials of service accounts responsible for AD-Entra ID synchronization and reset passwords by hybrid accounts to gain cloud access.
- Steal the NTLM hash of AZUREADSSOACC$ account: Forge tokens to authenticate as synced Entra ID users and take advantage of conditional access policies to bypass MFA.
- Federated Domain Manipulation: Switch a managed Entra ID domain to a federated one, and attackers establish a trust with a malicious domain, and can forge tokens to impersonate any user, bypassing MFA.

To accomplish eWective Active Directory and Entra ID attack surface hardening, an organization must reduce the attack surface by uncovering these and other security vulnerabilities and identity threats, mapping them against a standard security framework such as MITRE ATT&CK (https://attack.mitre.org/).

Actionable mitigation tasks to harden against actual attack techniques that threat actors have used in past attacks are recommended in the MITRE ATT&CK framework. Complete as many of those actions as possible to minimize attacks and remove vulnerabilities in your identify environment.

Best Practice Strategies and Actionable Steps

Specific actions you take for attack surface hardening or reduction in Active Directory (AD) and Entra ID are essential to safeguarding your organization's identity infrastructure. The following strategies and actionable steps will minimize vulnerabilities and limit opportunities for attackers.

The MITRE ATT&CK Framework helps in this regard. It suggests specific actions that can help detect, and respond to, the various threats the following strategies reference. For convenience and insights, references to MITRE ATT&CK mitigation recommendations are included with the strategies discussed here. Read more at https://attack.mitre.org/.

The strategies are:

- Implement Least Privilege Models
- Setup and Enforce Authentication and Access Controls
- Adopt Tiered Administration and Network Segmentation
- Harden and Monitor Service Accounts
- Apply Secure Configuration and Patch Management
- Mitigate Known Attack Techniques
- Continuous Monitoring and Auditing
- Encrypt and Backup Critical Data

Implement Least Privilege Models

a. Role-based permissions. Assign users the minimum permissions necessary for their roles. Avoid granting excessive privileges, especially for day-to-day tasks. For example,

- **Admin rights for specific tasks**

Use just-in-time or privileged access management (PAM) solutions to grant admin rights only when needed and revoke them automatically after the task is complete.

And instead of granting broad admin rights to allow specific tasks to be done, delegate only the necessary permissions for those tasks using custom groups or role-based access control (RBAC).

For example: the ability to add machine/computer accounts must be controlled closely, and the best practice is to delegate on the Create Computer objects permission to the specific OU where computer accounts should be created.

- **Access to sensitive resources**

For sensitive resources, use RBAC to systematically restrict access.

b. Review often. Regularly review and adjust group memberships and role assignments to ensure only authorized personnel retain elevated privileges.

c. Protected Groups. Place highly privileged accounts in protected groups, such as the "Protected Users" group, to enforce stricter security controls.

MITRE ATT&CK Mitigation for specific attack techniques to strengthen enforcement of least privilege model:

T1003 (OS Credential Dumping): Block standard users from local admin groups.
T1558 (Kerberoasting): Restrict service accounts to required SPNs only.
T1134 (Token Manipulation): Deny Standard Users.

Setup and Enforce Authentication and Access Controls

a. AS-REP Roasting. Enable Kerberos pre-authentication for all accounts to defend against AS-REP Roasting attacks. For any accounts that must bypass this setting, ensure strong passwords.

b. Strong passwords. Enforce strong, unique passwords (ideally passphrases of at least 30 characters) and consider using password protection features to block known breached passwords.

c. Monitor and correct. Monitor for and remediate accounts configured with insecure settings. Alert on those without Kerberos pre-authentication or with weak passwords, for example, to raise immediate awareness for corrective action.

d. Multi-factor authentication (MFA). Require MFA for all users, neutralizing threats from stolen passwords being used to infiltrate Active Directory or Entra ID and reducing the risk of lateral movement by attackers trying to compromise other accounts or resources in the Active Directory or Entra ID environment.

e. Entra ID security defaults. Enable security defaults in Entra ID to automatically enforce baseline security measures like MFA and blocking legacy authentication.

f. Zero-trust model. Require robust authentication for every connection. Removing implicit trust will make it much harder for attackers to succeed. And it strengthens defense against insider threats by either enforcing additional authentication, limiting permissions, or denying access completely.

g. Guest users. Limit guest user permissions and restrict who can invite external users to your directory.

h. IP allowlisting. Use IP allowlisting for network-level control to restrict access to trusted network locations. It's a network-level control that safeguards against external threats by unauthorized users accessing sensitive resources even if they have valid/stolen credentials. Since stolen credentials cannot be used from allowlisted IPs, phishing and credential stuffing are defended. For even stronger defense, use it along with OpenID Connect, which requires both valid credentials and a trusted IP for access – but trusted users will not need to enter credentials repeatedly once on an allowlisted network.

i. Stale account reviews. Through continuous or regular monitoring, review and promptly disable or delete inactive, orphaned, or unnecessary user and service accounts. This thwarts threats from insiders and former employees and eliminates an easy entry point for malicious actors.

j. Do not allow shared admin accounts. Review all existing synched Active Directory admin accounts in Entra ID, too, monitoring and auditing for suspicious behavior.

MITRE ATT&CK Mitigations of specific attack techniques to strengthen authentication and access controls:

T1556.001 (DC Authentication): Mitigate by enforcing MFA for domain controllers.

T1621 (MFA Spamming): Use time-bound MFA codes and alert on bulk requests.

T1040 (Network Sniffing): Encrypt authentication traffic and use challenge-response mechanisms.

Adopt Tiered Administration and Network Segmentation

a. Isolated Admin accounts. Separate administrative activities by implementing a tiered model (e.g., Tier 0 for domain controllers, Tier 1 for server admins, Tier 2 for end users). Admin accounts are then limited to a tier and cannot access higher tiers. This protects critical resources like domain controllers from a malicious actor pivoting from a compromised workstation, etc. Attackers cannot steal credentials for Tier 0 assets from lower tiers, so attack techniques like Kerberoasting and pass-the-hash are disrupted.

b. Isolated Admin workstations. Restrict administrative access to dedicated, hardened workstations that are isolated from internet access and general user environments.

c. Isolated Tier 0 systems. Domain controllers, Entra ID privileged roles, and any other core identity infrastructure components should be isolated from all other workstations and servers, minimizing risks of broader or crippling attacks.

d. Monitor and analyze traffic. Alert on cross-zone authentication attempts, (e.g., Tier 1 to Tier 0).

e. Penetration Testing. Simulate lateral movement attempts to validate segmentation.

MITRE ATT&CK Mitigation for specific attack techniques to strengthen tiered administration and network segmentation:

T1482 (Domain Trust Discovery): Isolate forest-domain boundaries.

T1536.002 (RDP Hijacking): Block RDP between security zones.

T1136 (Create Account): Restrict access to domain controllers.

T1098 (Account Manipulation): Segment identity management systems

Harden and Monitor Service Accounts

Service accounts are high-value targets for attacks in Active Directory and Entra ID because of their elevated privileges and broad access. They deserve special attention in your attempts to reduce the attack surface and repel attackers.

a. Group Managed Service Accounts (gMSAs). Wherever possible, use gMSAs in AD, or Managed Identities in Entra ID, for automated password management and rotation, reducing the risk of password-based attacks on service accounts.

b. Service Principal Names (SPNs). In Entra ID, limit the privileges of accounts with SPNs and avoid assigning them to high-privilege groups. Use them for non-Azure hosted services or multi-tenant applications.

c. Restrict access to hosts. Use the “Log On To” property to restrict whichcomputers a service account can access. Limit the scope of access to only necessary hosts.

d. Do not copy old service accounts. Create service accounts from scratch, to avoid inadvertent grants of excessive privileges.

e. Service Account Activity audits and monitoring. Regularly audit service account permissions and usage to detect and remediate unnecessary access.

f. AD Group Policy and OUs. Place service accounts in dedicated OUs and apply Group Policies to facilitate monitoring and enforce restrictions consistently.

g. Entra ID Conditional Access Policies. Restrict access based on device compliance, location, or risk.

MITRE ATT&CK Mitigation ID for privileged account management: M102

Apply Secure Configuration and Patch Management

Patch management is often considered the most effective activity for reducing the attack surface in AD, so reducing risk. Here is why:

a. Patch Active Directory regularly. Regularly update and patch all systems, including domain controllers and administrative hosts, to address known vulnerabilities, reducing ways in for attackers. Patching also proactively updates your defenses as known threats evolve, and new threats appear.

b. Server Core installations. Deploy domain controllers using the Server Core installation option, minimizing installed components on the server, so reducing possible vulnerabilities.

c. Reduce services and features. Disable legacy protocols (NTLM, SMBv1) and unused services (IIS, Print Spooler) on domain controllers.

d. Advanced auditing. Enable advanced auditing for Active Directory object changes (e.g., 'Audit Directory Service Changes'). Critical events are:

- Event 4738: User account changed.
- Event 4742: Computer account modified.

e. Organizational security policy reviews. Review and amend default security settings in Active Directory and Entra ID to constantly improve the organization's security posture by focusing on reducing the number of potential entry points for attackers. Activities should include:

- Inventory all Active Directory components, including DCs, privileged accounts, service accounts, and network connections.
- Map dependencies and interactions between Active Directory and other critical systems. Then, prioritize security controls and monitoring for those dependencies. And study the dependencies identified for unnecessary dependencies that could be exploited – directly reducing attack vectors in your environment.
- Apply the Tier Model accordingly in Active Directory to minimize risks of widespread impact from a single compromise.
- Fully Patch Entra ID hybrid components, like Entra Connect, and admin workstations.
- Take all the actions recommended in this section of the white paper.

MITRE ATT&CK Mitigation ID for secure configuration and patch management: M1054
Mitigations for related attack techniques:

T1543 (Create/Modify System Process): Enforce rootless containers for ADintegrated apps.

T1562.010 (Downgrade Attack): Disable legacy protocols and enforce encryption.

T1197 (BITS Jobs): Block BITS transfers on domain controllers.

Continuous Monitoring and Auditing

Another set of activities widely considered essential to improving the Active Directory and Entra ID security posture, monitoring to detect suspicious activity and threats quickly, and auditing, will help you detect potential threats early and investigate them to take appropriate action.

a. Monitor changes to objects. Monitor all changes to users, roles, groups, and policies. This includes tracking who made a change, when, and what was changed. Real-time alerts for critical changes (like role assignments or risky sign-ins) will help you respond quickly to threats.

b. Monitor activities. Monitor directory changes, privileged account activities, and authentication attempts to detect suspicious behavior quickly.

c. Pay close attention to delegations. Delegation of permissions can open doors to attackers, so audit delegations and monitor for excessive permissions. An example: Resource Based Constrained Delegation (RBCD) introduces critical security risks that can allow attackers to gain unauthorized access by delegating impersonation rights to malicious accounts. And lateral movement opportunities can open for the attacker.

d. Mitigate threats by auditing. Enable and configure auditing in Active Directory by using Group Policy Management on your domain controllers. In Entra ID, configure auditing similarly using either: 1) The Azure Portal and navigating to Monitoring, then select Audit Logs; or 2) Microsoft Defender or Purview to start recording user and admin activities from the Audit section.

Both systems' native auditing provides flexibility in what is captured, and both allow you to view logs.

Or use a third-party tool that captures events as they happen without the use of native Active Directory auditing. Audit successful account management, successful directory service access, and successful and failed logon events.

1. Mitigate threats by auditing relevant events such as configuration changes, group memberships, and privilege assignments. Maintain the resulting logs to help guide incident response and for compliance purposes.
2. For example, delegations configured as "unconstrained" or non-admin accounts with Machine Account Creation privileges should be deemed suspicious until proven otherwise. So,
    - Filter your auditing for activities like changes to configurations and for specific suspicious settings.
    - Ensure that the forensic trail of these and related actions and outcomes are stored securely to help with detection, investigation, and responses.

MITRE ATT&CK Mitigation of specific attack techniques through monitoring:

T1558 (Kerberoasting). Monitor detection rule: Spikes in occurrences of service tickets from a single account.

T1003.006 (DCSync). Monitor for unauthorized replication requests from non-domain controllers.

T1098 (Account Manipulation). Monitor for unexpected additions to privileged groups.

MITRE ATT&CK Mitigation of specific attack tactics through auditing:

T1003.006. Audit ACLs for replication permissions to prevent unauthorized delegation.

T1558.003. Audit Active Directory periodically for service accounts with Service Principal

Names (SPNs), ensuring their passwords and encryption types meet security standards.

T1134 (Access token manipulation). Audit command-line activity and API calls related to token manipulation.

T1558.001 (Steal or Forge Kerberos Tickets: Golden Ticket). Audit privileged accounts for unusual Kerberos activity.

Encrypt and Backup Critical Data

Security is constructed in layers. Encryption and data backups add critical layers to your overall security posture. Do not leave these out of your overall security plans.

Your last line of defense around sensitive and business-critical data is data encryption, making the **data unreadable** if an attacker gains access to your environment. For you, as a security practitioner, the data includes credentials, group policy objects, and configuration files. **Regular data backups** ensure business resilience if data are compromised by malicious actors.

Active Directory and Entra ID each have some unique encryption and backup considerations.

And MITRE ATT&CK Guidelines for encryption and data backups are identified below. Mitigation Mappings list attack techniques that intruders use and point you to actionable steps you can take to help you reduce the attack surface systematically while aligning to industry-standard practices.

Encryption

Encrypt sensitive data. Data at rest or in transit calls for protection against unauthorized access by going beyond Active Directory default settings to require strong encryption (e.g., AES-256) for authentication and data storage. Combined with continuous monitoring and regular audits, allowing early detection of suspicious activity and misconfigurations, the defenses of encryption are bolstered even more.

> ⇒ *Entra ID note: Use Microsoft's server-side encryption tools for data at rest, and ensure data in transit between apps, devices, and cloud services is protected by the likes of VPN or Transport Layer Security (TLS).*

Protect encryption keys. Securely manage encryption keys and restrict access to key management systems.

> ⇒ *Entra ID note: Secure encryption keys using solutions like Entra ID Key Vault and restrict access using IAM controls.*

Strengthen Authentication Protocols: Use AES encryption for Kerberos tickets instead of vulnerable algorithms like RC4 to mitigate Kerberoasting (T1558.003) and AS-REP Roasting (T1558.004)

- Secure private keys with hardware modules (TPM/HSM) to prevent credential theft.

MITRE ATT&CK Mitigation Mapping for Encryption.
Mitigation ID: **M1041**
Key Techniques Addressed (with detection tips):

T1557 (Adversary-in-the-Middle)

T1558 (Steal or Forge Kerberos Tickets)

T1003 (OS Credential Dumping)

T1040 (Network Sniffing).

T1530 (Data from Cloud Storage)

Backups

Prepare for business continuity

Encrypt backup files. Store them in encrypted locations, both onsite and offsite, to prevent attackers from gaining access to sensitive Active Directory data.

Backup domain controllers at least twice per domain for redundancy and resilience. During a ransomware attack, it will be possible to restore data without paying a ransom, if reliable backups are available.

Immutable Backups. Use write-once-read-many (WORM) storage to prevent ransomware from altering or deleting backups.

Offsite and Isolated Storage. Store backups in cloud services (e.g., AWS S3, Azure Backup) or physical offsite locations to ensure availability during onsite breaches.

Regular Testing: Validate backup integrity by periodically restoring them to confirm recoverability.

Secure Backup Systems. Harden backup servers and networks to prevent compromise during incidents like Data Manipulation (Technique T1556, MITRE ATT&CK).

MITRE ATT&CK Mitigation Mapping for Backups.
Mitigation ID: **M1053**
Key Techniques Addressed:

T1486 (Data Encrypted for Impact)

T1490 (Inhibit System Recovery)

By systematically applying these strategies and taking these actions, organizations can significantly reduce the attack surface of both Active Directory and Entra ID, improving their overall security posture and resilience against identity-based threats.

Conclusion

As a security practitioner, you are among those responsible for studying and defending against cyberattacks. You must help lead the effort to address vulnerabilities and make infiltration by attackers as difficult as possible.

The spread of exploits within your organization depends on you and other security pros taking actions described in this paper to reduce the attack surface.

Reducing the Active Directory and Entra ID attack surface will not guarantee that all attacks will stop. But doing your due diligence to put in place safeguards under your control will be your best defense against disruptive or disastrous cybercriminal attacks.

Cygna Labs Corp

sales@cygnalabs.com
Toll Free: 844.442.9462
Intl: +1 (305) 501-2430
www.cygnalabs.com