semperis

Essential Guide to Securing Microsoft Active Directory

WHITE PAPER HOW TO UNCOVER SECURITY VULNERABILITIES IN YOUR CORE IDENTITY SYSTEM

semperis

Do You Know Your Active Directory Security Vulnerabilities?

By Sean Deuby
Semperis Director of Services

Securing Microsoft Active Directory (AD) involves dealing with a mixed bag of risks, ranging from management mistakes to unpatched vulnerabilities. AD has become a prime target for cyber-attackers who use AD to elevate privileges and gain persistence in the organization. Investigate a typical data breach, and you'll find that stolen credentials likely were used—sometimes for initial entry, sometimes for accessing critical systems, but always to the detriment of the targeted organization.

Hardening AD begins with getting a handle on the vulnerabilities and common configuration and management mishaps that pave the road to compromises. To defend AD, administrators need to know how attackers are targeting their environment. How many, however, can pass a pop quiz about the types of security holes threat actors are sneaking through as they move through the steps of the breach?

Authentication fail

Secure password policies should be the order of the day throughout the Active Directory infrastructure. It's important to note these long-held password policies are changing, but when they change, they should change as a set. Specifically, user password expiration and rotation policies have been shown to weaken a password's strength because humans fall into easily predictable patterns when forced to update regularly. However, these policies should not be changed until common passwords have been eliminated from the directory (perhaps using Azure AD Password Protection). Any account with the PASSWD_NOTREQD flag set should automatically draw additional scrutiny and have a justifiable reason for its configuration.

Additionally, passwords—especially service account passwords—should be periodically rotated. Leaving passwords unchanged for lengthy amounts of time increases the likelihood of a successful brute force attack, as attackers will have more time to take swipes at them.

Authentication issues to watch for include:

1. Computers and Group Managed Service Accounts (gMSA) objects with passwords set over 90 days ago
2. Reversible passwords found in Group Policy Objects (GPOs)
3. Anonymous access to Active Directory enabled
4. Zerologon vulnerability (CVE-2020-1472) if the patch is not applied.

"Weak passwords, non-expiring passwords, no passwords—all these are warning signs that an organization's AD environment is not secure."

Permitting excessive permissions

As most AD environments have been in production for many years, their attack surfaces have grown.Many of a forest's accumulated vulnerabilities can be traced back to the pattern that someone needs something done, usually in a hurry, and the least–privilege path to get that done is too time consuming, not easily available, or simply not known. As a result, the user or group or permission is over-privileged just to ensure the request will be satisfied and the ticket closed. And of course, that entitlement is never ever removed, so the attack surface simply grows and grows.

In reality, it's not uncommon for AD environments to have unnecessarily high numbers of domain administrators—a fact that can be even more troubling if those accounts are orphaned and are simply waiting to be leveraged in an attack. Service accounts with excess permissions also pose a high risk because their passwords are usually set to not expire, and many of them will have weak passwords (which makes them a good kerberoasting target). As the number of users with administrative privileges grows, so does the attack surface that needs to be protected. Membership to these groups should be tightly controlled.

As the number of users with administrative privileges grows, so does the attack surface that needs to be protected.

Mistakes happen, of course. As an AD environment grows larger and more complex, for example, someone might fail to properly account for inherited permissions and inadvertently grant an account too many privileges. But even properly managing privilege delegation is not enough with attackers taking the offensive.

As an example, consider the impact of an AdminSDHolder attack. Just as a refresher, the AdminSDHolder container stores the Security Descriptor applied to privileged groups. By default, every 60 minutes, the Security Description Propagation (SDPROP) process compares the permissions on protected objects and reverses any discrepancies according to what is defined in AdminSDHolder.

In an AdminSDHolder attack, threat actors exploit SDPROP to maintain persistence by replacing the permissions of an object with the attacker's unauthorized modifications. If the permission changes are identified and undone, but the unauthorized changes to AdminSDHolder are undetected, the attacker's changes will be reinstated.

Auditing permissions and monitoring for suspicious activity is the best defense against the abuse of privileges.

Permission issues to watch for include:

1. Privileged objects with unprivileged owners
2. Permission changes on the AdminSDHolder object
3. Unprivileged users with DC Sync rights on the domain
4. Default security descriptor schema changes in the last 90 days

> *"Purple Knight addresses a need that has become more pronounced in the wake of the Exchange Server Hafnium attack, which prompted Microsoft to advise customers to scan their systems for IOEs and IOCs. Large, complex organizations tend to have a spider web of permissions that have accumulated over time—and no idea whether that situation can be exploited."*
>
> Darren Mar-Elia, Semperis VP of Products

Cheat sheet for security

Armed with information about indicators of exposure (IOEs), organizations can strengthen their AD's security. One tool that can help is Purple Knight, a free security assessment tool from Semperis. Purple Knight queries your Active Directory environment in "read-only" mode and performs a comprehensive set of tests against the most common and effective attack vectors to uncover risky configurations and security weaknesses.

Scanning Active Directory provides insight into its security posture and reduces the risk of unauthorized changes or misconfigurations going undetected. AD administrators need to know more than their craft; they also need to know the tactics of their adversaries. By keeping critical warning signs top of mind, they can harden AD against common attacks.

FREE SECURITY TOOL DOWNLOAD

Spot weaknesses in Active Directory before attackers do.

Active Directory holds the "keys to the kingdom," and if not safeguarded properly, it will compromise your entire security infrastructure. Purple Knight is a free Active Directory security assessment tool built and managed by an elite group of Microsoft identity experts.

Request access to Purple Knight →

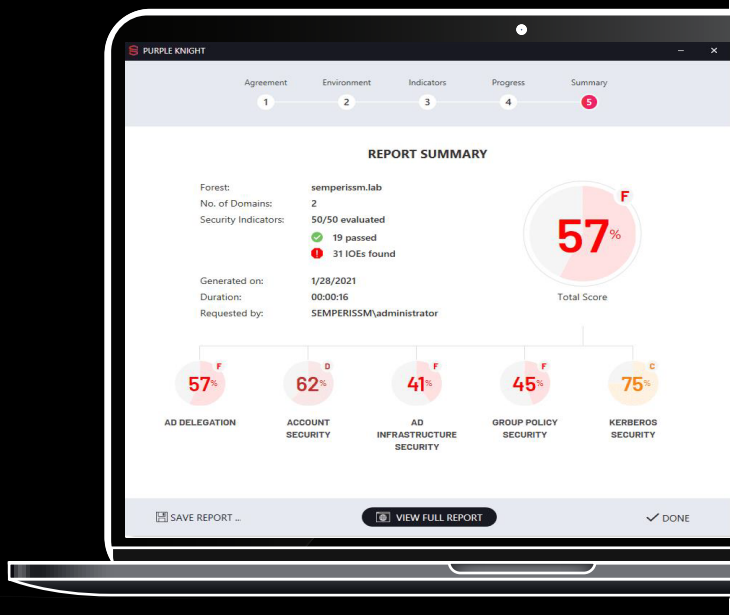

semperis

How to Defend Against Active Directory Attacks That Leave No Trace

By Guido Grillenmeier
Semperis Chief Technologist

Cybercriminals are using new tactics and techniques to gain access to Active Directory in novel ways, making their attacks even more dangerous—and more necessary to detect. One of the most important parts of any cybersecurity strategy is detection. Having an ability to spot the bad guy entering, moving about, or worse—administering—your network is key to a swift response. And with the median number of days an attacker sits undetected on your network at 146, according to Microsoft, it's evident that the bad guys are very good at working in stealth.

When it comes to detecting potentially malicious actions within Active Directory (AD), most organizations rely on Domain Controller event log consolidation and SIEM solutions to spot abnormal logons and changes. This all works—as long as the attack technique leaves a log trail. A few types of attacks have been seen in the wild that leave no discernible trail or, at least, any evidence of malicious activity. Some examples include:

- **DCShadow attack:** Using the DCShadow functionality within the hacker tool Mimikatz, this attack first takes the path of registering a rogue domain controller (DC) by modifying the Configuration partition of AD. Then the threat actor makes malicious fake changes (e.g., changes to group memberships of Domain Admins, or even less obvious changes such as adding the SID of the Domain Admins group to the sidHistory attribute of a compromised normal user). This attack technique bypasses traditional SIEM-based logging, as the rogue DC doesn't report the changes. Instead, changes are injected directly into the replication stream of the production domain controllers.
- **Group Policy change:** A documented attack involving Ryuk ransomware resulted in changes being made to a Group Policy object that propagated the installation of Ryuk to remote endpoints within the victim organization. By default, event logs don't include details on what was changed within a Group Policy. So, if an attacker makes a malicious change (as in the case of Ryuk), all that's seen is that an account with access to the Group Policy made a change, which probably won't set off any alarms.
- **Zerologon attack:** After a proof–of–concept exploit code was released in public, an attacker with network access to a domain controller was able to send special Netlogon messages consisting of strings of zeros, forcing the domain controller computer password to be changed to an empty string. So, without any logon—i.e., with zero logon—the attacker now owns the domain controller, can perform any changes in AD, and can further use this path to attack other systems in your infrastructure. It is unlikely that your monitoring tools today are watching out for unexpected password changes on your DCs.

"A few types of attacks have been seen in the wild that leave no discernable trail or, at least, any evidence of malicious activity."

It isn't by chance that these attacks don't leave a trace; it's by design. The bad guys are spending massive amounts of time inspecting exactly how their target environments function and looking for ways to bypass, obfuscate, and circumvent any form of detection—which includes logging.

Because these kinds of attacks exist, the question becomes what should you do about it—both proactively and reactively?

Protecting against malicious Active Directory changes

There are three ways to protect your organization against malicious AD changes:

1. **Monitor AD for malicious changes:** This goes beyond SIEM and involves a third-party solution designed to see every change made within AD—regardless of who makes it, on which DC, using what solution, etc.—ideally by reading and understanding the replication traffic of the DCs themselves. This monitoring needs to include changes within Group Policy as well. In many cases, solutions designed to monitor changes in AD can define specific protected objects to be monitored for any change—for example, changes in membership to Domain Admins—so that anytime those protected objects are modified, alarms do go off. The solution should cover both changes to Group Policies as well as visibility into replication.

2. **Look for DCShadow:** Mimikatz leaves some artifacts behind and there are some telltale signs that DCShadow has been used on your network. Reviewing AD for these signs needs to be part of a regular review of AD security. Note that once you find a trace of Mimikatz DCShadow in your environment, you must act quickly as you'll already be a victim of an attack. At that point, you will wish you also had a solution that would show you what changes were performed at the replication level, which you could then analyze and ideally revert.

3. **Be able to recover AD**: Your organization needs the proactive ability to recover any and all of AD should you determine that AD has been compromised. In some cases, you can be thinking in terms of backups and a DR strategy to recover AD in a cyberattack scenario. Should you indeed need to recover your complete AD service, potentially as the next victim of a malware attack, beware that a good domain controller backup does not equate to a seamless and fast AD service recovery. You'll want to have practiced the whole recovery process periodically, following the copious Microsoft AD Forest Recovery Guide. But it's equally valuable to look for solutions that can revert changes down to the attribute level or even automatically revert changes to protect objects when detected.

Targeting Active Directory and modifying it to suit the attacker is a common tactic taken by today's cybercriminal—so much so that the old model of watching AD audit events for changes might no longer be viable. Organizations that are serious about the security and integrity of their AD need to be looking for additional ways to gain visibility into every AD change and have the ability to revert or recover when necessary.

+1-703-918-4884
info@semperis.com
www.semperis.com

221 River Street
9th floor
Hoboken, NJ 07030

Semperis is the pioneer of identity-driven cyber resilience for cross-cloud and hybrid environments. The company provides cyber preparedness, incident response, and disaster recovery solutions for enterprise directory services—the keys to the kingdom. Semperis' patented technology for Microsoft Active Directory protects over 40 million identities from cyberattacks, data breaches, and operational errors. Semperis is headquartered in New York City and operates internationally, with its research and development team distributed between San Francisco and Tel Aviv.

Semperis hosts the award-winning Hybrid Identity Protection conference. The company has received the highest level of industry accolades; most recently being named Best Business Continuity / Disaster Recovery Solution by SC Magazine's 2020 Trust Awards. Semperis is accredited by Microsoft and recognized by Gartner.