CSAI Foundation | Cloud Security Alliance

Certighost: AD CS Flaw Lets Users Impersonate Domain Controllers

2026-07-28

AI-assisted Rapid Research

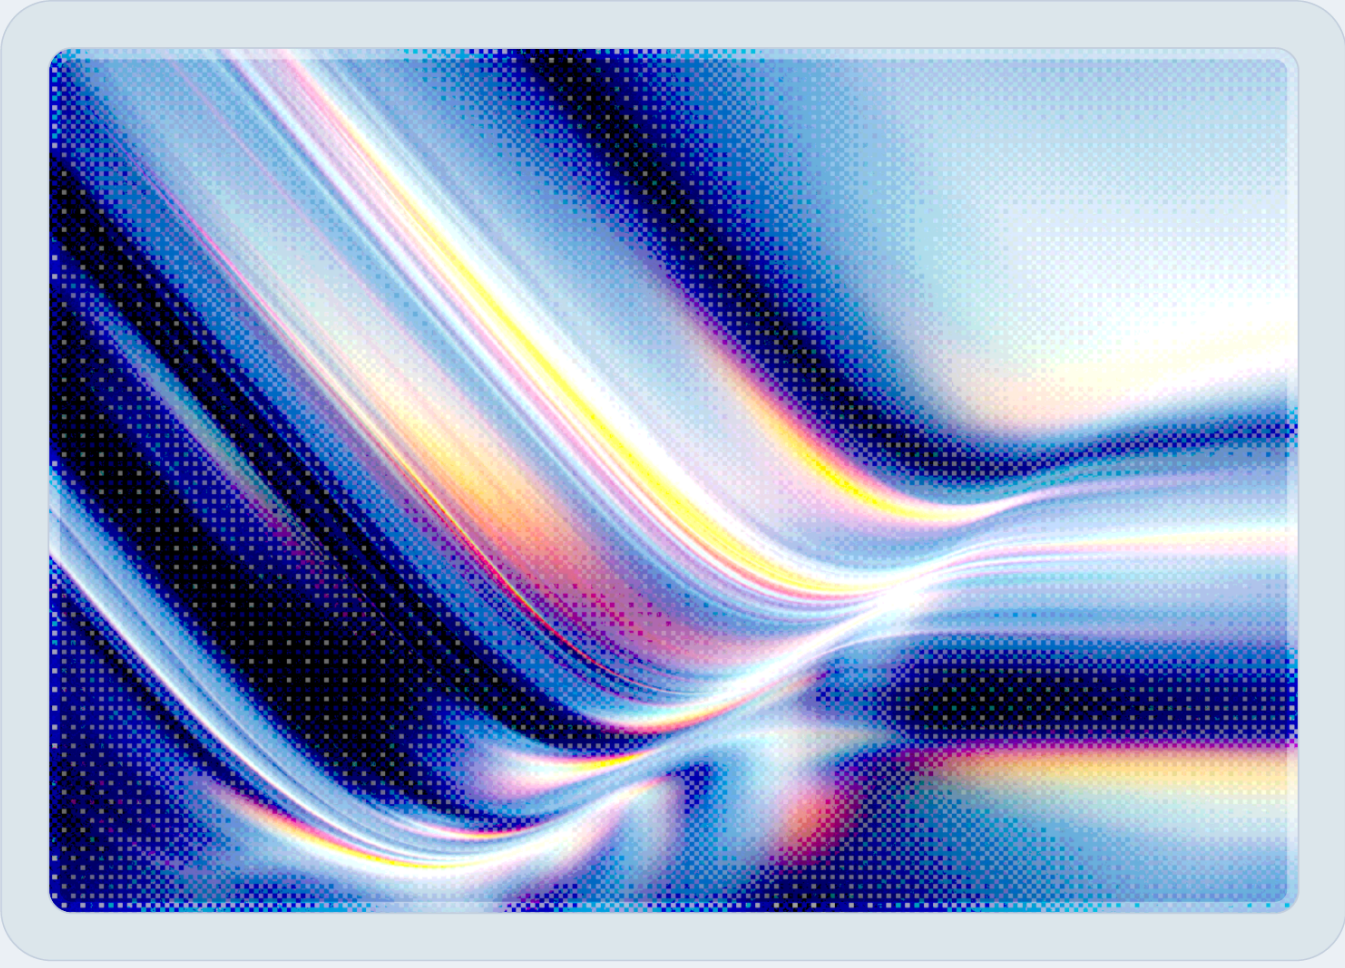

CSAI

CSA cloud security alliance®

Key Takeaways

A newly disclosed vulnerability in Microsoft Active Directory Certificate Services, tracked as CVE-2026-54121 and nicknamed Certighost, allows any authenticated domain user without administrative rights to obtain a certificate that lets them authenticate to Kerberos as a domain controller [1][2]. Researchers H0j3n and Aniq Fakhrul privately reported the flaw to Microsoft on May 14, 2026; Microsoft confirmed it on May 22 and shipped a fix in its July 14, 2026 security updates, and the researchers released technical detail and a working proof-of-concept exploit on July 24, 2026 [1]. The flaw carries a CVSS score of 8.8 and affects Windows Server 2012 through Server 2025, including Server Core builds, as well as Windows 10 versions 1607 and 1809 [1]. Exploitation requires only a standard domain account and network reachability to a vulnerable enterprise certification authority, with no administrator privileges or user interaction needed [3][4]. Because the resulting forged domain controller credential carries directory replication rights, an attacker can perform a DCSync operation to extract the krbtgt secret and, from there, compromise the entire Active Directory forest [1][2]. Organizations that operate Active Directory Certificate Services should treat the July 2026 patch as urgent, verify it has been applied, and hunt for signs of prior abuse, since public proof-of-concept code now exists and Microsoft had initially assessed exploitation as "less likely" before the researchers' disclosure changed that calculus [4].

Background

Active Directory Certificate Services has for years been a favored target for attackers pursuing privilege escalation and lateral movement inside Windows domains, because a certification authority sits at a uniquely trusted junction: it issues credentials that Kerberos will treat as proof of identity for any account whose attributes the CA is willing to vouch for. Prior research into this attack surface, most notably the "Certified Pre-Owned" family of misconfigurations published by SpecterOps in 2021, showed how permissive certificate templates or CA enrollment policies could let low-privileged users mint certificates that impersonate arbitrarily powerful accounts, including domain controllers themselves. Microsoft has spent several years hardening default template permissions and adding certificate mapping controls in response to that research, so many security teams that had already remediated the classic ESC1-ESC8 misconfigurations likely assumed their AD CS deployments were in a materially safer position by 2026.

Certighost demonstrates that the underlying trust model still contains gaps beyond template permissions. Rather than abusing a misconfigured certificate template, the vulnerability lives in a lesser-known enrollment fallback mechanism that Windows calls a "chase," used when a certification authority needs to resolve directory information about the entity requesting a certificate [2][3][5]. The chase mechanism accepts two client-supplied request attributes, commonly referred to as cdc (identifying the domain controller the CA should contact) and rmd (identifying the machine object to resolve), and prior to the July 2026 patch, the CA did not verify that the host named in the cdc attribute was actually a legitimate domain controller [2][5]. That gap turned a piece of enrollment plumbing into a pathway for an attacker to redirect the CA's identity-resolution queries toward infrastructure the attacker controls.

The researchers' write-up ties the flaw to default Active Directory behavior that is often left unchanged in production: the ms-DS-MachineAccountQuota attribute, which by default lets any authenticated domain user register up to ten new computer accounts [1][3]. That default, combined with the unauthenticated trust the CA previously placed in the cdc value, is what makes Certighost exploitable from a baseline "any domain user" starting position rather than requiring an attacker to already hold elevated rights or a foothold on a privileged host.

Security Analysis

The exploitation chain begins with an attacker who holds nothing more than a standard domain account. Using the default machine account quota, the attacker registers a computer account, or reuses one already available, and then submits a certificate enrollment request to a vulnerable enterprise CA against the default Machine certificate template [1][3]. Within that request, the attacker sets the cdc attribute to point at attacker-controlled infrastructure rather than a real domain controller, and sets the rmd attribute to reference the target: a domain controller's computer account. When the CA attempts to chase and resolve that target's identity information, it reaches out to the host named in cdc, which the attacker has configured to run rogue LSA, LDAP, and SMB services [1][2][5].

At that point the attacker's rogue services relay the CA's authentication attempt to a real domain controller over Netlogon, effectively bridging two protocol conversations so that the CA believes it has successfully validated the target's identity when it has in fact been fed relayed and manipulated data [1]. The CA, satisfied by this relay, issues a certificate binding the identity of the targeted domain controller to a public key the attacker controls. Because the certificate is technically valid and correctly chained to the CA, the attacker can then present it during Kerberos PKINIT authentication and be issued a ticket-granting ticket as if they were the domain controller itself [1][2]. Domain controller accounts inherently carry the Replicating Directory Changes and Replicating Directory Changes All permissions needed for Active Directory replication, so a Kerberos identity spoofed to that level of privilege can invoke DCSync

and pull password hashes for any account in the domain, including the krbtgt account whose key underlies the security of every Kerberos ticket issued in that forest [1][2][3]. An attacker holding the krbtgt secret can subsequently forge Golden Tickets, persist across password resets, and move freely across any system that trusts the domain.

What distinguishes Certighost from earlier AD CS escalation paths is that it does not depend on an administrator having misconfigured a certificate template's enrollment permissions or issuance policy; it exploits a validation gap in code that ships with default AD CS behavior, which means organizations that had already hardened their templates against the well-known ESC-series misconfigurations were not necessarily protected [2][4]. That characteristic elevates the vulnerability's practical severity, since remediation cannot be achieved through template permission review alone and instead requires the vendor patch or an equivalent configuration change to the chase mechanism itself.

Detection guidance published alongside the disclosure recommends enabling Certificate Services auditing and watching for anomalous certificate-issuance events, specifically Windows event IDs 4886 and 4887, which record certificate request and issuance activity on the CA, along with any Microsoft Defender alerts referencing potential Certighost or CVE-2026-54121 abuse patterns [1]. Because the attack chain depends on relayed Netlogon authentication traffic reaching a rogue service, network monitoring for CA hosts initiating unexpected outbound authentication attempts to non-domain-controller hosts is also a reasonable detection heuristic. CSA's review of public sources as of this writing found no published Sigma rule or comprehensive indicator set specific to this technique.

Recommendations

Immediate Actions

Organizations running Active Directory Certificate Services should confirm that the July 14, 2026 Microsoft security update has been applied to every certification authority in their environment, since this patch adds validation that requires the cdc-supplied hostname to resolve to a genuine domain controller computer object, with a matching DNS name, the SERVER_TRUST_ACCOUNT flag set, and a SID comparison to block object substitution [1][5]. Given that a public proof-of-concept now exists, unpatched CAs should be treated as an active, exploitable exposure rather than a theoretical risk, and patch verification should take priority over routine change-management timelines. Where immediate patching is not possible, Microsoft's disclosed workaround disables the chase flag on the CA policy module with certutil -setreg policy\EditFlags -EDITF_ENABLECHASECLIENTDC followed by a restart of

the Certificate Services service, though the researchers and multiple outlets caution that this workaround has not been exhaustively tested against all enrollment scenarios and should be treated as a bridge measure rather than a substitute for patching [1][2].

Short-Term Mitigations

Security teams should enable detailed Certificate Services auditing if it is not already active, and add monitoring for event IDs 4886 and 4887 on every CA, correlating unusual certificate issuance patterns against the Machine template with the identity of the requesting account [1]. Reviewing and tightening the ms-DS-MachineAccountQuota attribute, which defaults to ten and is rarely revisited after domain provisioning, reduces the ease with which a low-privileged account can stage the computer accounts this attack chain depends on. Teams should also review existing certificate issuance logs retroactively for anomalous requests referencing domain controller accounts from non-administrative user contexts, since the technique's public availability means opportunistic exploitation attempts are likely to follow disclosure.

Strategic Considerations

Certighost is a reminder that AD CS hardening efforts focused solely on template permissions and enrollment agent restrictions, the areas emphasized by the well-known ESC1 through ESC8 misconfiguration classes, do not cover every path by which a certification authority's trust decisions can be subverted. Organizations should treat AD CS as a Tier 0 asset subject to the same change control, monitoring, and least-privilege scrutiny applied to domain controllers themselves, and should periodically reassess enrollment protocol behaviors, not just template permissions, as Microsoft continues to patch this attack surface. Longer term, organizations pursuing Zero Trust architectures should ensure that identity issuance infrastructure, including certification authorities, is included within continuous verification and monitoring scope rather than treated as a one-time hardening exercise, since PKI trust anchors that are compromised can undermine identity assurances made everywhere else in the environment.

CSA Resource Alignment

Certighost is fundamentally an identity spoofing and identity abuse problem: an attacker with a low-privileged identity manipulates a trust decision made by supporting infrastructure to obtain a credential that Kerberos accepts as belonging to a far more privileged identity, the domain controller itself. CSA's "Using Zero Trust to Counter Identity Spoofing & Abuse" [6] directly addresses this class of problem,

providing a taxonomy of how malicious actors subvert identity ecosystems through spoofing and abuse of legitimate identity mechanisms, along with Zero Trust-aligned detection and mitigation strategies. Certighost's exploitation of a CA's implicit trust in a client-supplied hostname during the chase process, and the subsequent forging of a domain-controller-level Kerberos identity, is a concrete real-world instance of the identity spoofing patterns that paper analyzes, and organizations applying its guidance on continuous identity verification would be better positioned to detect the anomalous authentication and certificate issuance activity this exploit generates.

More broadly, CSA's Zero Trust program supplies the architectural foundation for reducing this class of risk. The "Zero Trust Guiding Principles" [7] establish that no component of an identity infrastructure, including a certification authority, should be implicitly trusted based on network location or protocol role alone, a principle directly relevant to the chase mechanism's prior, unauthenticated trust in an attacker-supplied server hostname. Organizations should also evaluate this vulnerability against the identity and access management domain of CSA's AI Controls Matrix (AICM), since as enterprises increasingly connect AI agents and automation to Active Directory-backed identity systems for authentication and authorization, a domain controller impersonation technique like Certighost represents a potential upstream compromise path that could undermine the trust assumptions those AI systems rely on for their own access decisions.

References

[1] The Hacker News. "Certighost Exploit Lets Low-Privileged Active Directory Users Impersonate a Domain Controller." The Hacker News, July 2026.

[2] BleepingComputer. "New Certighost PoC exploit lets attackers hijack Windows domains." BleepingComputer, July 2026.

[3] Cybersecurity News. "Certighost Active Directory CS Exploit Allows Low-Privileged Users to Compromise Domain." Cybersecurity News, July 2026.

[4] Help Net Security. "PoC exploit released for critical AD CS domain-takeover flaw (CVE-2026-54121)." Help Net Security, July 27, 2026.

[5] H0j3n. "CVE-2026-54121 Technical Write-up." GitHub Gist, July 2026.

[6] Cloud Security Alliance. "Using Zero Trust to Counter Identity Spoofing & Abuse." Cloud Security Alliance, 2026.

[7] Cloud Security Alliance. "Zero Trust Guiding Principles." Cloud Security Alliance.

