ultraviolet

THREAT ADVISORY

VibeCoded Malware Targeting Active Directory

Services Performed By:
UltraViolet Cyber TIDE Team
tide@uvcyber.com

Published Date:
July 15, 2026
TLP:GREEN

Executive Snapshot

Threat actors are now using generative AI to write bespoke, single-use PowerShell scripts for Active Directory enumeration. Huntress recently documented this technique following a June 2026 intrusion where an attacker used pre-compromised credentials and RDP access to deploy an AI-generated script that mapped users, computers, groups, and trusts before exfiltrating the data via legitimate tools like s5cmd. The script's AI origin was evident from leftover prompt artifacts, unedited placeholder text, redundant fallback logic, and decorative console output, none of which affected its core function. A parallel case from Sygnia showed the same dynamic playing out in cloud environments, where an AI-assisted attacker compromised a large AWS environment in about 72 hours using only known techniques executed at unprecedented speed. The real risk to organizations is not a new attack method but accelerated, harder-to-fingerprint execution of familiar tradecraft, since each AI-generated script is functionally unique and evades traditional hash and signature-based detection.

- Enforce multifactor authentication and just-in-time access controls on privileged accounts and RDP-facing systems to reduce the impact of compromised credentials.
- Monitor and restrict common staging directories such as C:\ProgramData, and flag unusual use of legitimate but exfiltration-capable tools like s5cmd.
- Update incident response plans and tabletop exercises to account for compressed attack timelines, since AI-assisted actors can execute reconnaissance and exfiltration far faster than traditional playbooks anticipate.

TIDE Team Analysis

Researchers at Huntress recently reconstructed and analyzed a PowerShell script, Untitled1.ps1, used by a threat actor during a June 3, 2026 intrusion against a domain-joined Windows Server. The attacker gained access via RDP using pre-compromised credentials, staged tools in C:\ProgramData, and within minutes deployed the script to enumerate the Active Directory environment. The script exhibited unmistakable hallmarks of AI generation: a title reading "100% Working AD Information Gathering Script - FULLY FIXED," an unedited placeholder hostname left over from the AI's example output, a redundant five-method fallback chain to locate the domain controller, and decorative, color-coded console output that serves no operational purpose. This represents one of the first well-documented, real-world cases of a threat actor using generative AI to author bespoke, single-use attack tooling rather than relying on established frameworks like BloodHound or Cobalt Strike.

The risk this trend poses to enterprise environments is less about novel attack techniques and more about speed, scale, and detection evasion. The underlying attack chain in this incident, credential compromise, RDP pivot, AD enumeration, and data exfiltration via legitimate tools like s5cmd and SharpShares, followed a familiar data theft pattern. What changed is that AI lowered the technical bar for less-skilled actors to produce functional, custom malware on demand, and because each generated script is unique, traditional hash-based and static signature detection struggle to flag it. A related case documented by Sygnia illustrates the same dynamic at cloud scale: an AI-assisted attacker moved from initial access to broad compromise of an AWS environment in roughly 72 hours by chaining known techniques (credential harvesting, discovery, persistence, and impact activities) far faster than defenders could respond, rather than by using any new exploit or zero-day.

For enterprises, this means the volume and velocity of attacks against identity infrastructure and cloud environments is likely to increase even without a corresponding rise in attacker sophistication. Active Directory remains an especially attractive target because full domain enumeration, the kind these scripts perform, gives an attacker the map needed to identify high-value accounts, shares, and trust relationships for follow-on privilege escalation or data theft. Organizations that depend heavily on signature-based antivirus or EDR as a primary control are at elevated risk, since these AI-authored scripts are functionally novel each time they are generated and will not match any known hash or string signature.

The core defensive takeaway from both incidents is that AI changes the code but not the underlying behavior of an attack. Enumerating a domain controller, dumping user and group data, and moving data to cloud storage all leave a consistent operational footprint regardless of how the script performing them was written. Organizations should prioritize behavioral detection and telemetry-based monitoring, particularly around PowerShell Script Block Logging (Event ID 4104), unusual LDAP or AD module queries, and anomalous use of legitimate but exfiltration-capable tools such as s5cmd. Credential hygiene and RDP exposure reduction remain foundational, since both incidents began with compromised credentials and direct or VPN-based RDP access rather than a technical exploit.

In practical terms, security leadership should treat this development as a reason to accelerate, not fundamentally redesign, existing security roadmaps. Recommended actions include enforcing multifactor authentication and just-in-time access for privileged and RDP-facing accounts, tightening monitoring on ProgramData and other common staging directories, deploying or tuning SIEM detections around AD enumeration behaviors rather than known tool signatures, and extending the same behavioral monitoring philosophy to cloud environments given the Sygnia findings on AI-accelerated cloud compromise. Tabletop exercises should be updated to reflect compressed attack timelines, since AI-assisted actors are demonstrating the ability to compress days or weeks of manual reconnaissance into minutes or hours. The threat is an acceleration of known tradecraft rather than a new category of attack, and organizations with mature identity security and behavioral detection programs are well positioned to adapt.

Why It Matters

For nearly two decades, enterprise defense has leaned heavily on identifying known offensive tooling: file hashes for Cobalt Strike, string signatures for PowerSploit, behavioral fingerprints for BloodHound. That model assumed attackers were largely reusing a shared, finite pool of frameworks, giving defenders a stable set of artifacts to detect against. Vibe coding breaks that assumption. When any threat actor, regardless of skill level, can iteratively prompt an AI into producing a working, custom enumeration script in minutes, the industry loses the reusable-artifact advantage that signature-based detection was built on. This is not a hypothetical shift; it has already been observed in a live intrusion, which suggests the practice is likely to spread quickly given how low the barrier to entry has become.

Looking forward, this trend is likely to compress the skills gap between low-tier and advanced threat actors, since AI tooling can substitute for years of scripting expertise. It also raises the operational tempo of attacks, as seen in the Sygnia AWS case, where AI-assisted actors executed a full compromise lifecycle in 72 hours using nothing but well-known techniques. As generative AI models continue to improve at code generation, defenders should expect an increasing volume of similarly bespoke tooling across other attack surfaces, including cloud infrastructure, endpoint malware, and phishing kits, not just Active Directory enumeration.

The industry-wide implication is a necessary pivot from artifact-based detection toward behavior-based detection as the primary line of defense. Security programs that continue to over-index on signatures, hashes, and known-bad indicators will see a growing detection gap as more attackers adopt AI-assisted tooling. Programs built around monitoring the fundamental mechanics of an attack, credential use, enumeration patterns, data movement, are better positioned to remain effective regardless of how the malware performing those actions was written or how often it changes.

How to Respond

- Strictly adhere to cybersecurity fundamentals and ensure all personnel undergo annual phishing and social engineering training. Speak with your UltraViolet Cyber TAM Representative to schedule a live phishing engagement.
- Review and restrict RDP exposure, including enforcing multifactor authentication on all remote access paths and eliminating direct internet-facing RDP where possible.
- Perform annual tech refresh reviews to gain a holistic understanding of your infrastructure. Speak with your UltraViolet Cyber TAM Representative to schedule a Red Team or Purple Team engagement to gain insight into the vulnerabilities in your environment.

What UltraViolet Cyber is Doing

- Tracking new CVEs and high impact vulnerabilities, analyzing and deploying public Proof-Of-Concept code against custom built targets.
- Proactively enabling custom detections based on the collected artifacts, tactics, techniques, and procedures identified in this activity.
- Performing hypothesis driven threat hunts based on threat actor behavior and artifacts. UVCyber customers will be informed of the results through secure channels.
- Parsing available victim dump data for any social, financial, business, or technical relations to UVCyber Clients and partner organizations.
- Aggregating threat intelligence from myriad sources and applying the most up-to-date knowledge to proactive threat hunting and response.

About UltraViolet Cyber

UltraViolet Cyber is a leading tech-enabled managed security services provider, delivering unparalleled cybersecurity expertise that fills technology and talent gaps across Global 2000 and Federal Government customers. Founded and operated by security practitioners from the national intelligence community, UltraViolet Cyber connects offensive security, application security, detection and response, and security engineering to deliver a differentiated approach to cybersecurity operations. Transforming customers' security programs, UltraViolet Cyber's flagship security-as-a-service solution, UV Lens, removes complex operational silos, replacing them with integrated security capabilities. UltraViolet is headquartered in McLean, Virginia with technology centers across the world.

443.351.7630 / info@uvcyber.com | UltraViolet Cyber | @uv_cyber

---