SPECTEROPS
Creators of BloodHound

Applying Attack Path Management to Identity Security

How Gaps in Your Identity Stack Leave Room for Attackers

Introduction

It's no secret that security is getting more difficult. Evolving threats like AI-enabled attackers, the proliferation of ransomware-as-a-service, and nation-state supported groups provide nearly seemingly infinite resources to adversaries. The continued accumulation of identity tech debt through rapid digital transformation efforts has expanded the attack surface significantly. Nonhuman identities have overtaken human identities in volume, further exacerbating the challenge of understanding threats to the enterprise.

Countless alerts have been added to the Security Operations Center to attempt to create a perimeter around the most critical assets in the environment—administrative accounts, customer information, regulatory and sensitive data, and security-minded responses to these perceived threats are implemented to disrupt or control attacks. **But adversaries are altering their techniques and procedures to avoid these solutions entirely, or they are using modern tactics** like purchasing stolen credentials or stealing authentication tokens in flight to appear as a normal user to these tools while advancing their operations without detection.

As adversary tactics change and organizations continue to focus on a perimeter that is quickly disintegrating, defenders and their leadership need to pause and ask themselves, "Are we winning?" The controls are in place and are doing what they are designed to do, but we need to focus on what these solutions aren't doing and start evaluating environments like an attacker does: **look for the same gaps and exposure threat actors seek and eliminate them pre-exploit, focusing on work that creates real progress.**

This problem was highlighted recently in a series of campaigns attributed to the ransomware group Storm-0501.

CASE STUDY

Storm-0501 Hybrid Attack Highlights Gap in Identity Security Stacks

Seeing the Anatomy of an Attack Through the Eye of the Storm

Starting in August 2024, Microsoft Threat Intelligence released detailed analyses of how the threat actor Storm-0501 evolved from executing traditional on-premises ransomware attacks to conducting highly sophisticated, cloud-based operations. The report illustrates a striking example of an adversary seamlessly transitioning between on-premises Active Directory (AD) and Azure tenant compromises, exploiting identity relationships and misconfigurations that connect the two environments.

What makes this campaign particularly alarming is how it weaved through a complex hybrid environment, bypassing many conventional security controls in the process. Despite widespread deployment of endpoint protection, identity monitoring, conditional access with MFA, and SIEM integrations, Storm-0501 was able to persist and laterally move across hybrid environments, highlighting the limitations of post-exploit defenses when complex and often hidden identity relationships themselves form the attack surface.

This paper examines how the same identity attack paths used by the Storm-0501 actors can be proactively identified, visualized, and mitigated through effective attack path management (APM).

By continuously mapping and analyzing the graph of identities, permissions, and trust relationships across AD and Azure/Entra ID, APM provides defenders with the visibility and control needed to prevent adversary movement before exploitation occurs.

The Storm-0501 campaign serves as a powerful case study in how understanding and managing attack paths programmatically transforms identity security from reactive detection to proactive defense.

The modern adversary doesn't need to create new tools for success, instead they exploit what already exists: legitimate identities, widespread misconfigurations, long-standing AD/identity technical debt, and the ability to acquire or manipulate valid credentials. The Storm-0501 campaigns demonstrate a pattern seen at SpecterOps repeatedly:

1. Initial access via exposed/unpatched services or stolen credentials
2. Follow-on discovery using native tools
3. Credential harvesting and secret extraction
4. Lateral movement using existing offensive security tooling[1]
5. Finally, a pivot from on-prem to cloud via Microsoft Entra Connect or compromised sync/service accounts, enabling tenant takeover and either backdoors or ransomware deployments

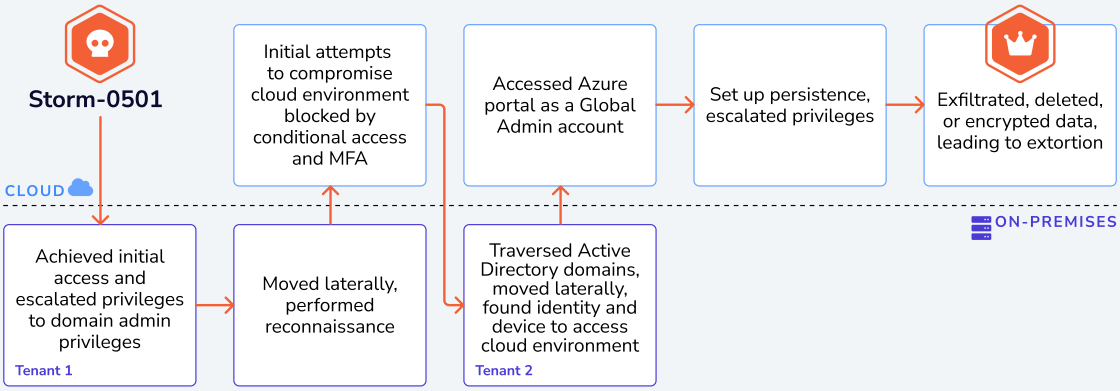

Storm-0501 used these methods to bypass perimeter defenses and traditional identity tools by acting as legitimate identities. By exploiting unpatched application vulnerabilities, the adversary obtained administrative control early in the intrusion. Once a foothold was established, they used native Windows tools and open-source utilities for discovery of valuable targets, leveraging legitimate administrative privileges to evade endpoint detection.

The detection and response solutions involved focused on behavioral anomalies and endpoint compromise that were ineffective due to the use of valid credentials and legitimate management tools, which then allowed the actors to blend in with normal, expected user behaviors and traffic.

Identity Governance (IGA) and Privileged Access Management (PAM) solutions could not have prevented lateral movement, as they only manage explicit entitlements and initial credential use, not post-authentication secrets abuse or the myriad of cascading relationships and misconfigurations that formed the real attack paths that were ultimately used in the attack.[2]

Through repeated credential theft and lateral movement, Storm-0501 reached Domain Admin privileges, exfiltrated sensitive data, and pivoted to Microsoft Entra ID by exploiting weaknesses in synchronization accounts and MFA policies. This identity snowball reaction ultimately led to cloud compromise and ransomware deployment.

As hybrid identity configurations become commonplace, it is essential to understand the impact of how attackers can leverage legacy identity tech debt along with the authorization drift that comes with increasingly cloud-based environments to move silently between these realms. **Traditional security controls address important aspects of this problem, but there are gaps in and between their coverage that create opportunities to advance the attack.**

1. microsoft.com/en-us/security/blog/2025/08/27/storm-0501s-evolving-techniques-lead-to-cloud-based-ransomware
2. specterops.io/blog/2025/10/08/the-clean-source-principle-and-the-future-of-identity-security

These tools are good at what they do.

But they have gaps that adversaries leverage.

The Gaps

EDR / ITDR (Endpoint Detection & Response / Identity Threat Detection & Response)

Detected post-compromise behavior but missed legitimate credential re-use and privilege escalation.

ITDR is strong at spotting threats at the system and behavioral level, such as identifying suspicious logins, privilege misuse, and changes to the directory that align with known attack signatures, but both try to operate after the abuse has started and don't address the identity tech debt that supply the attack paths adversaries use to evade detection, create persistence, and identify targets for privilege escalation.

ITDR and EDR miss the underlying identity graph, the attack paths used by adversaries to move laterally, and the use of what appear to be entirely legitimate actions to evade detection as they further their attack.

PAM (Privileged Access Management)

Controls access to privileged credentials at rest but is ineffective after access is granted.

PAM solutions are great at controlling initial credential use (multi-factor authentication, just-in-time, conditional access, etc.), but after the credential is used, session artifacts linger, where they may be accessed by adversaries.

Additionally, coverage gaps in most PAM implementations exist due to an inability to identify all the hidden, incidental access that is missed during the initial discovery and rollout of the PAM solution. Finding explicitly granted privilege in the obvious places is straightforward (check for user accounts, service accounts, application users, and administrative accounts that require privilege). This excludes objects like GPOs, Resource Groups, and OUs that can be leveraged by an attacker to gain control over privileged identities. And, what about accounts with abusable permissions enabling them to move laterally or escalate privileges? This risk is missed because it is extremely difficult to discover.

These gaps are particularly dangerous because they represent risk that is not obvious or known but are just as dangerous as the inherent privilege they easily discover and include.

IGA (Identity Governance & Administration)

Described who should have access but not the implicit, cascading privileges that let an attacker become someone else.

Most IGA tools provide a clear picture of who can access or log into what. This data is not difficult to gather and provides a starting point for ongoing assessment of access needs. But this view paints an incomplete picture of risk. Attackers typically don't use their initial access victim to reach their objective. They typically need to pivot/laterally move to another identity to assume more privilege in the process. In other words, attackers don't care what their current victim has access to now; they care about who they can become.

IGA is great for provisioning access at scale and informing you who should have access to certain resources but does not account for cascading privileges or how control of one identity could lead into control over another or multiple other identities. Attackers know this and target these gaps in coverage to evade IGA coverage, creating persistence, and opportunities to move throughout the organization—one step closer to privilege.

PKI (Public Key Infrastructure) Misconfigurations

Allowed certificate abuse for undetectable persistence.

Many organizations leverage PKI within their identity platform to create a private, internal system for issuing and managing digital certificates used to provide secure authentication to network resources. This is commonly achieved by using Microsoft Active Directory Certificate Services (ADCS) to automate the deployment and management of these certificates and improve the security of access to critical internal assets.

However, despite the capabilities offered by ADCS, **it has security shortcomings that can actually degrade the security of an organization**, as was discovered by members of the SpecterOps research team.[3] Attackers can find ADCS servers, and any misconfigured access templates or broad enrollment permissions and access certificates with elevated privileges, such as domain administrator certificates. Once a certificate is obtained, an attacker can move throughout the environment completely undetected, because they have become the target account, and establish persistence to gain access to critical assets and applications.

CNAPP / CIEM (Cloud Native Application Protection Platform / Cloud Infrastructure Entitlement Management)

Provided partial cloud visibility but failed to connect on-prem identity risks to cloud compromise.

Manages and monitors user permissions and access privileges, mitigating the risk of data breaches from excessive or inappropriate access.

CNAPP technology tends to fail because they rely on extensive and highly complicated integrations that can be difficult to align with existing security tooling. CNAPP is intended to provide a holistic view of cloud security, but wide-ranging and complex integrations are required to create this ecosphere. When not integrated seamlessly across dynamic cloud environments, coverage gaps emerge, making prioritization of risk difficult—resulting in false positives or misunderstood alerts, and potentially delaying responsive actions. Both can have a dampening effect on actual risk mitigation.

Further, CNAPP solutions lack insight into how cascading control can allow one identity to take over other identities or resources under their control until they reach the privilege or access they need. **Attackers look at identities in these terms—not just what can be done with an account, but what other accounts can be owned and leveraged by that account.**

3. specterops.io/blog/2021/06/17/certified-pre-owned

BloodHound Enterprise fills those gaps.

Attack Path Management Gives Control Back to the Security Team

Attackers evolve with security; it's as simple as that. Merely adding another hurdle to protect your vital assets isn't enough. A fundamental and foundational understanding of identities and their intended and unintended relationships, and how those relationships can be connected to create direct pathways to those protected assets is needed.

Security culture should switch from reacting to breaches to preventing attacks. BloodHound Enterprise's Attack Path Management program gives you this visibility and empowerment to programmatically eliminate attack paths and their underlying risk in a strategic, prioritized process that focuses on maximizing risk reduction with every action.

BloodHound Enterprise builds a comprehensive attack graph of identities and resources across platforms to produce a comprehensive attack map to answer the question, "If one of my identities or resources are breached, what can happen to the larger enterprise?"

BloodHound Enterprise models the identity graph and its actionable edges, finds the attack paths (who can become whom) with the most impact, and prescribes targeted remediations that break those paths without disrupting business. **Where the other tools detect behavior or enumerate entitlements, BloodHound Enterprise reduces the structural risk within your environment**.

An Attack Path Management program, such as that delivered through BloodHound Enterprise, can:

- **Map and prioritize hidden attack paths** created by over-privileged identities, misconfigured permissions, and hybrid trust or cross-platform relationships before exploitation
- **Systematically eliminate lateral movement opportunities** for attackers, rather than monitor for their execution with ITDR / EDR
- **Identify exploitable connections** between on-prem and cloud identities
- **Quantify identity risk** across human and non-human accounts, reducing privilege sprawl that facilitated persistence
- **Provided actionable, programmatic remediation guidance**, ensuring continuous identity hardening beyond static compliance checklists

Implementing an Identity Attack Path Management Program

30/60/90 Roadmap (Practical, Prioritized, Manageable)

0–30 days (rapid reduction)

- Run BloodHound Enterprise to establish a baseline Attack Graph for the environment across on-premises and cloud systems
- Use the data to identify unused privileged identities
- Identify and remediate attack paths to Entra Connect hosts; rotate sync credentials; require tenant-level MFA for Global Admins
- Immediately remediate the top 5 highest-impact attack paths (as identified and prioritized by BloodHound Enterprise)
- Harden PKI: lock certificate templates, test certificate issuance monitoring

30–60 days (structural hardening)

- Move shadow administrative accounts under PAM (JIT + session recording) to strengthen PAS coverage
- Ensure RMM server and service accounts are included in protection by Privilege Zones if they also have the ability to manage or deploy to other critical systems, including the Entra Sync server

60-90 days (programmatic prevention)

- Integrate APM into change control: every new identity, group, policy, etc. change is assessed and tested with BloodHound Enterprise to determine impact
- Automate the APM remediation cycles to measure and report path closure

Key metrics to measure effectiveness

- **Total number of all identities and resources monitored and number of privileged identities and resources under protection**
- **Number of unique attack paths to critical identities and resources (before vs after)**
- **Number of attack paths remediated (ongoing)**
- **Mean Time to Remediate (MTTR) for attack paths**
- **Number of Entra Connect/ ADCS configuration changes blocked or put through high-assurance workflow**

Storm-0501's campaign succeeded because it did not rely on traditional exploits and bugs but leveraged legitimate identities and legacy configuration debt that traditional security controls simply could not see. This is another indication that adversaries are evolving their tactics, techniques, and procedures to avoid modern security tooling and take advantage of these gaps in coverage.

You can have best-in-class EDR, CNAPP, PAM, IGA with CAP/MFA and systems and still be vulnerable if the paths between identities, especially on-prem based identities, and resources remain intact.

BloodHound Enterprise reveals those paths, quantifies their risk, and prescribes safe, precise fixes, so you are continuously solidifying your baseline identity foundation in a risk-focused and manageable process. Identity Attack Path Management turns siloed security investments into layered security and gives defenders the ability to stop adversaries from using legitimate tools in a takeover.