SPECTEROPS

ATTACK PATH MANAGEMENT

Maturity Model

Executive Summary

The rise of identity-based attacks has exposed a blind spot in the security stack; none of the core tools were designed to map or control how privilege chains together. Attackers aren't slipping through gaps in detection; they're bypassing controls entirely by chaining together legitimate access into attack paths that lead directly to critical assets.

Identity Attack Path Management (Identity APM) is a new practice within security to address chains of abusable privileges and user behaviors that create direct and indirect connections between computers and users. As attackers abuse these attack paths, it's clear that reactive detection and surface-level access reviews are not enough. What's needed is a practice that treats privilege relationships as the control surface, and manages them accordingly. But as with any emerging capability, organizations are struggling to benchmark where they are, what good looks like, and how to improve.

This white paper introduces the Identity Attack Path Management Maturity Model, a structured framework to evaluate how effectively an organization identifies, prioritizes, and eliminates attack paths. The model defines six levels of maturity and is based on the Capability Maturity Model Integration (CMMI) model created by Carnegie Mellon University. Each level is detailed across the dimensions of People, Process, and Technology, allowing organizations to identify not just tooling gaps, but the organizational and procedural blockers holding them back.

This paper is written for security, identity, and infrastructure leaders who recognize that lateral movement and privilege escalation are not just detection problems; they are structural failures of control. Most organizations have no reliable way to measure how effectively they're enforcing least privilege across hybrid identity systems. Identity APM fills that gap, offering both visibility and a quantifiable framework to assess exposure. By adopting this model, teams can baseline their current maturity, identify key investments needed to advance, and align Identity APM as a continuous, measurable practice within their broader security strategy.

What Is Identity Attack Path Management?

Identity APM is the continuous discovery, mapping, and risk assessment of attack path choke points; the critical privilege relationships and structural exposures that enable lateral movement and privilege escalation.

Key characteristics:

Choke-point oriented: Prioritizes key relationships and privilege escalations that cause the widest blast radius

Cross-platform: Covers hybrid infrastructure; on-prem, cloud, and everything in between

Graph-driven: Models identities, groups, roles, sessions, tokens, etc. combine into effective privilege

Continuous: An ongoing risk-reduction cycle rather than a periodic audit

By analysing identity structure, Identity APM shifts organisations from reactive cleanup toward structural control.

Identity APM Maturity Levels at a Glance

● **Level 0 – Nonexistent**
Reactive only—attack paths are unknown; response only occurs after incidents or pen tests. No visibility, ownership, or tracking.

● **Level 1 – Initial**
Ad hoc—remediation is manual, sporadic, and untracked; ownership is informal. No measurement or prioritization exists.

● **Level 2 – Managed**
Repeatable—documented processes run on a schedule, yet efforts remain siloed and qualitative. Risk prioritization is subjective.

● **Level 3 – Defined**
Enterprise tooling is in place; attack path remediation is continuous and risk-informed with alignment across Security, Identity, and Infrastructure. Organizations begin using metrics to assess path volume, criticality, and coverage.

● **Level 4 – Quantitatively Managed**
Identity APM insights are embedded in incident response, provisioning, and governance workflows. Common KPIs and a shared risk model drive response across functions. Teams trust Identity APM data as the source of truth for identity privilege exposure.

● **Level 5 – Optimizing**
Attack paths are avoided at design time. Identity APM data actively shapes provisioning templates, access decisions, and architectural patterns. Preventive controls are enforced via policy, and metrics track not just risk reduction, but path prevention.

Why a Maturity Model for Identity APM?

Security teams have long relied on maturity models to evaluate how well they perform in key areas: vulnerability management, incident response, identity governance, and beyond. Without a view of privilege chains, those initiatives burn time chasing symptoms instead of addressing the root cause. But no such framework has existed for managing attack paths, even as they become the dominant mechanism for lateral movement and privilege escalation in modern environments.

Most organizations are flying blind. Some have started remediating attack paths manually. Others have adopted tools like BloodHound Enterprise but tool output needs owners who track closure, and raw findings need governance and metrics before they become a language executives can act on. Without a shared vocabulary, leadership cannot gauge posture or justify spend.

The Identity Attack Path Management Maturity Model fills that gap.

Identity APM is a structural, collaborative practice. When SecOps opens a ticket, IAM can see the same graph and close the loop faster; vulnerability analysts can filter CVEs that do not intersect a live attack path, cutting noise and focusing patch windows. Identity APM focuses on how identities and privileges chain together across on-prem and cloud systems to form exploitable paths; paths that traditional access control tools miss entirely. Unlike ITDR which triggers on behavior, Identity APM removes the attack path before the behavior can take place. These paths often stem from nested group memberships, delegated administration, hybrid misalignments, or overlooked relationships between accounts and roles.

What makes Identity APM unique is that it treats effective privilege and identity coercion, not just assigned privilege, as the true source of risk. It forces organizations to confront not just what access was granted, but what access was created through complex identity relationships. And it brings structure and visibility to an area that has traditionally been driven by spreadsheets, assumptions, or incident retrospectives.

The Identity APM Maturity Model introduces a structured way to:

- **Baseline your current state**—understand what level of capability your organization actually has.
- **Identify specific gaps**—not just in tooling, but in ownership, process rigor, and integration.
- **Measure Identity Risk**—monitor and track risk-reduction over time.
- **Chart a path forward**—provide a roadmap for how to evolve Identity APM from reactive cleanup to proactive design.

This model is grounded in real-world Identity APM deployments[1] and lessons learned across enterprises at various stages of maturity.

It is designed to be practical, diagnostic, and prescriptive, giving organizations a way to move beyond ad hoc efforts and establish Identity APM as a formal, measurable security function.

[1] specterops.io/wp-content/uploads/sites/3/2023/11/Sustaining-and-Continuous-Verification-of-Tier-0-with-BloodHound-Enterprise.pdf

Understanding the Scope of the Problem

In nearly every major breach, attackers aren't breaking in, they're logging in and moving laterally through access that already exist. Here are three recent examples:

MGM[2]

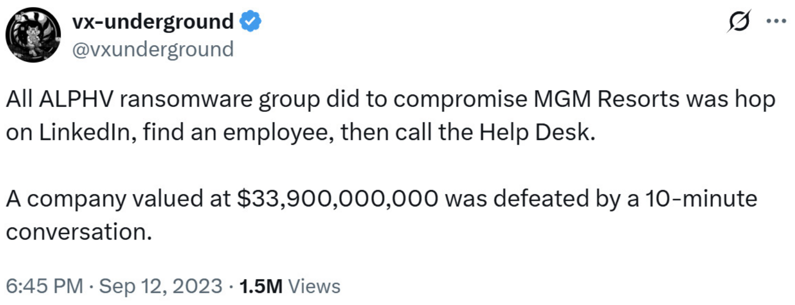

vx-underground
@vxunderground

All ALPHV ransomware group did to compromise MGM Resorts was hop on LinkedIn, find an employee, then call the Help Desk.

A company valued at $33,900,000,000 was defeated by a 10-minute conversation.

6:45 PM · Sep 12, 2023 · **1.5M** Views

Okta[3]

"The unauthorized access to Okta's customer support system leveraged a service account stored in the system itself. This service account was granted permissions to view and update customer support cases. During our investigation into suspicious use of this account, Okta Security identified that an employee had signed-in to their personal Google profile on the Chrome browser of their Okta-managed laptop. The username and password of the service account had been saved into the employee's personal Google account. The most likely avenue for exposure of this credential is the compromise of the employee's personal Google account or personal device."

Marks and Spencer[4]

"The compromise began with a simple but highly effective tactic: impersonation. Threat actors posing as an M&S employee contacted a third-party provider and convinced them to reset access credentials. The vendor, unaware of the deception, granted the request, unknowingly handing over a critical entry point. This allowed the attackers to escalate privileges and move laterally across connected systems."

[2] https://www.forbes.com/sites/suzannerowankelleher/2023/09/13/ransomware-attack-mgm-resorts
[3] https://sec.okta.com/articles/2023/11/unauthorized-access-oktas-support-case-management-system-root-cause
[4] https://conosco.com/in-the-news/marks-and-spencer-what-it-means

The fundamental issue isn't malware or misconfiguration. It's structure: the complex, often invisible ways identities accumulate privileges and interconnect across an environment.

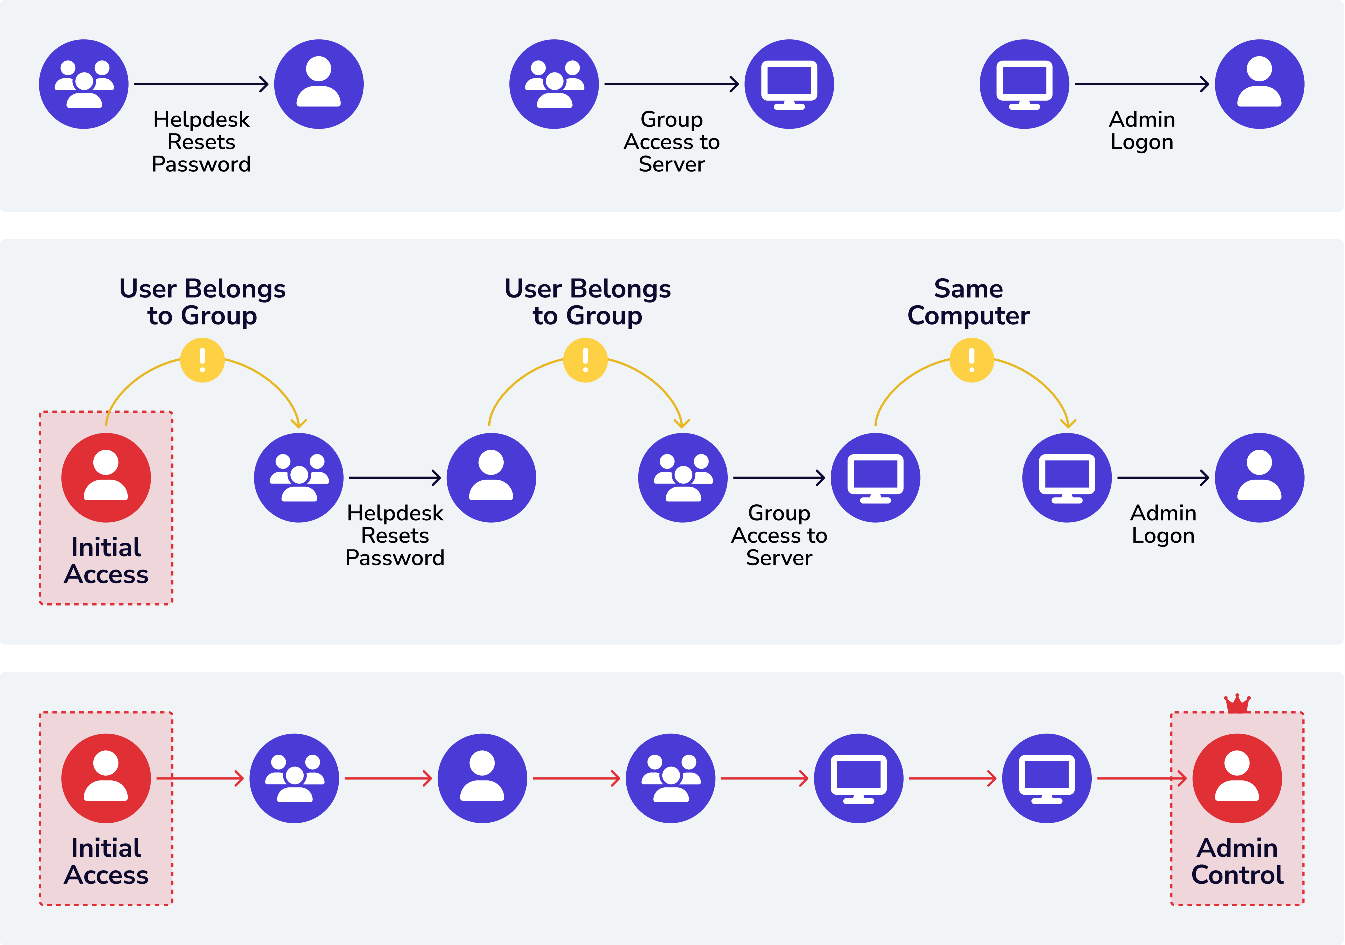

Attack paths are chains of exploitable relationships—nested group memberships, stale sessions, delegated admin rights, default roles—that connect ordinary user identities to critical resources and assets. These paths are often legitimate by the company's standards and even partners or external auditors. However, when viewed through the lens of an attacker, they form high-risk privilege escalations and lateral movement opportunities.

This problem isn't theoretical. In modern hybrid environments:

- **Thousands of identities** can generate millions of potential paths. A large enterprise with tens of thousands of identities can generate well over ten million potential paths once effective privileges are mapped.
- **Tools like IGA and PAM** see the surface-level assigned entitlements, but miss the inherited, effective access across domains and platforms that attackers exploit.
- **EDR and ITDR** detect behavior after it happens; by the time an alert fires, the lateral move is usually complete.

What's made the problem worse is the explosive growth and diversity of systems, identities, and interconnections in today's enterprise. Hybrid environments introduce new privilege relationships spanning from legacy Active Directory environments to Azure Entra ID, PaaS, SaaS platforms, cloud IAM layers, and workloads.

Non-human identities (NHIs)—like service accounts, automation bots, and workload identities—further increase the attack surface, often with less oversight and more privilege than human users.

Additionally, commodity AI tooling has lowered the barrier to exploitation. What once required expert-level analysis or red team tradecraft can now be automated, customized, and scaled. AI will also expand the use of Non-Human Identities which creates even more attack paths.

Even More Identities	Enhanced Attackers
Every AI agent needs an identity	Automated, continuous, agentic
NHIs outnumber humans 20-to-1	Nation-state tradecraft, for anyone
150% growth in NHIs this year	Too fast. Too many. Too late.

Despite these trends, most organizations are still operating without any formal program for managing attack paths. They lack the visibility, ownership, and repeatable processes needed to reduce this form of risk; let alone prevent it. Even those with strong security tooling in place often underestimate how easily and quickly an intruder can traverse their environment. The result: a massive disconnect between perceived security and actual exposure.

Not All "Attack Paths" Are Equal

Visualising a path is not the same as governing it. Many CSPM or cloud-IAM dashboards draw diagrams, yet they typically:

- Focus on **static cloud infrastructure**.
- Miss pivots that **cross trust boundaries** (for example, on-prem to cloud),
- Ignore **transitive privilege** or **nested relationships**.
- Don't take into account attacker techniques that can be used to expand available attack paths.

Identity Attack Path Management closes these gaps by analysing the full relationship graph—user to group, role to session, token to admin—and ranking every route by the impact it can deliver.

The emphasis shifts from where access resides to where it leads, giving defenders a control surface they can measure and reduce.

Introducing the Identity APM Maturity Model

Identity APM is neither a product nor a one-off project. It is an emerging security practice that demands coordination across people and processes, backed by the right technology. As Threat and Vulnerability Management once did, Identity APM now needs a maturity model so organisations can benchmark capability, align teams, and plan progression.

The **Identity APM Maturity Model** defines six levels of capability, ranging from complete absence of practice to preventive privilege design. What makes this model distinct is its multidimensional structure: each level is described through the lenses of **People, Process, and Technology**; the foundational elements of operational maturity.

This structure helps organizations measure far more than tool adoption; it shows whether roles, workflows, and integrations are in place to make Identity APM effective at scale.

Key Characteristics of the Model

- **Level 0 – Nonexistent:** Reactive only—attack paths are unknown; response only occurs after incidents or pen tests. No visibility, ownership, or tracking.
- **Level 1 – Initial:** Ad hoc—Remediation is manual, sporadic, and untracked; ownership is informal. No measurement or prioritization exists.
- **Level 2 – Managed:** Repeatable—Documented processes run on a schedule, yet efforts remain siloed and qualitative. Risk prioritization is subjective.
- **Level 3 – Defined:** Enterprise tooling is in place; attack path remediation is continuous and risk-informed with alignment across Security, Identity, and Infrastructure. Organizations begin using metrics to assess path volume, criticality, and coverage.
- **Level 4 – Quantitatively Managed:** Identity APM insights are embedded in incident response, provisioning, and governance workflows. Common KPIs and a shared risk model drive response across functions. Teams trust Identity APM data as the source of truth for identity privilege exposure.
- **Level 5 – Optimizing:** Attack paths are avoided at design time. Identity APM data actively shapes provisioning templates, access decisions, and architectural patterns. Preventive controls are enforced via policy, and metrics track not just risk reduction, but path prevention.

By mapping maturity across these levels, organizations gain a diagnostic view of where they stand and what's holding them back. For some, the blocker is tooling. For others, it's lack of process or ownership. The model surfaces those gaps and guides targeted improvement.

How to Use This Model

This maturity model is designed to be **diagnostic**, **prescriptive**, and **repeatable**. Use it as a structured tool to:

- **Assess Your Current State:** Start by identifying where your organization sits across the **People**, **Process**, and **Technology** dimensions. Be honest; many organizations find themselves straddling two levels. That's normal.
- **Pinpoint Your Gaps:** Look for asymmetry. Are you at Level 3 in tooling but still Level 1 in ownership or process? This model helps you spot where progress is blocked and why Identity APM efforts might be stalling.
- **Prioritize Next Steps:** Each level implies specific investments; new roles, automation, playbook development, or architectural changes. Use the model to build your roadmap and justify those moves to stakeholders.
- **Track Progress Over Time:** Revisit the model quarterly or annually. Treat it like you would a vulnerability management maturity assessment or incident response readiness check.

Detailed Level Breakdown

The Identity APM Maturity Model defines six levels, from zero capability to fully preventive design. Each level is assessed across **People**, **Process**, and **Technology** ensuring organizations evaluate not just tools, but the human and operational factors that make Identity APM effective in practice.

0 Level 0 – Nonexistent

At this stage, the organization has no awareness of attack paths as a security concern and actions are reactionary. Any exposure is discovered incidentally, typically through external assessments or after an intrusion. There is no ownership, no process, and no tooling in place.

- **People:** No ownership. Attack path risk is not recognized by Security, Identity, or IT teams.
- **Process:** No processes. Remediation is a reactive process post-incident or during audit outbriefs. No ongoing tracking.
- **Technology:** None. Organizations may receive CSVs or PDFs from red teams, but no structured data or tooling is used.

1 Level 1 – Initial

Here, an individual or small team has begun to understand attack paths and tries to remediate them ad hoc and manually. Efforts are inconsistent, undocumented, and driven by local context or intuition, not formal risk analysis.

- **People:** Individual practitioners or red teamers take initiative. No formal ownership, responsibility, or support.
- **Process:** Manual, sporadic cleanup efforts. No schedule, tracking, or prioritization.
- **Technology:** Community tools (e.g., BloodHound CE). No integration or automation.

2 Level 2 – Managed

The organization has recognized attack path management as a recurring concern. Teams begin to formalize their repeatable efforts, introducing documentation, ownership, and periodic reviews. Efforts primarily driven from one part of the organization (typically security) with sporadic support from other departments like identity and IT. Partial or pilot deployment of APM tooling, often scoped to a single domain like cloud IAM or infrastructure. These tools may surface local paths but fail to expose cross-domain risk, hybrid privilege escalation, or non-obvious identity relationships. Risk analysis remains largely subjective.

- **People:** Defined roles begin to emerge within a specific department. A small team tracks efforts but limited interaction across organizational lines (IT, Identity, etc.).
- **Process:** Internal processes documented. Quarterly or semi-regular reviews established.
- **Technology:** Pilot or limited deployment of Identity APM tooling. Still lacks full coverage or integration.

3 Level 3 – Defined

Identity APM is treated as a formal enterprise practice. Organizations have tooling in place, continuous remediation workflows, and risk-informed prioritization. Critically, this is where alignment across Security, Identity, and Infrastructure begins to emerge. Shared understanding of attack path data enables more effective decision-making, and organizations begin to converge on a common risk language and metrics around privilege. Attack paths are actively eliminated across environments. Not just after incidents, but as part of regular operations.

- **People:** Dedicated team or function for Identity APM. Cross-functional coordination with Security, IT, and Identity is streamlined and consistent.
- **Process:** Continuous monitoring and remediation. Attack paths considered during change management and provisioning. Metrics in place to track attack path volume, criticality, and coverage.
- **Technology:** Enterprise Identity APM tooling (e.g., BloodHound Enterprise) deployed across key environments, with quantifiable risk scoring.

4 Level 4 – Quantitatively Managed

At this level, Identity Attack Path Management is no longer siloed. It's embedded into the broader operational fabric of the organization, with attack path data actively informing investigations, access decisions, incident response, and change management. Teams don't just remediate, they collaborate.

Alignment becomes a defining feature. Security, Identity, and Infrastructure teams are now operating from a shared understanding of risk, informed by a common graph and consistent definitions of effective privilege. Attack path data becomes a trusted, cross-functional source of truth—not just a tool for one team.

This alignment enables faster remediation, clearer prioritization, and more effective incident response. It also sets the foundation for shifting from cleanup to prevention.

- **People:** Shared accountability across Security, Identity, and Infrastructure. Identity APM data and insights are routinely used across teams.
- **Process:** Attack path visibility is built into investigations, provisioning checks, and change reviews. Teams coordinate around shared KPIs and risk definitions.
- **Technology:** APM tooling is integrated into SOAR, SIEM, and ticketing systems. Graph data and choke point insights drive automated workflows and collaborative decision-making.

Cross-Functional Alignment: The Real Differentiator

The most mature Identity APM programs don't just have better tools; they have better alignment.

Leading organizations:

- Use a **shared graph of privilege** to unify understanding across Security, Identity, and Infrastructure.
- Align on **common KPIs** (e.g., path count, choke point density, blast radius).
- Build **joint response playbooks** that turn Identity APM insights into action across teams.
- Converge on a **shared language of risk** and move past tribal definitions of "admin," "privileged," or "sensitive."

These organizations don't just reduce paths faster—they reduce friction between the teams responsible for keeping privilege under control.

5 Level 5 – Optimizing

Organizations at this level don't just fix attack paths; they prevent them from forming. Identity APM data is used to inform provisioning templates, privilege design, and architectural decisions. Guardrails enforce least privilege by default. This is full-spectrum prevention, not post-facto cleanup.

- **People:** Identity APM expertise embedded into security architecture, identity governance, and engineering design teams.
- **Process:** Preventive controls drive access decisions. Identity APM metrics shape access reviews, template development, and organizational policy. Metrics move beyond risk reduction and now track path prevention.
- **Technology:** Automated enforcement via policy-driven controls (e.g., Privilege Zones). Attack paths are blocked before they're even created.

Level	People	Process	Technology
0 Nonexistent	No ownership or awareness of attack paths	No process; reactive cleanup only	None or pen test artifacts only
1 Initial	Individual efforts only; no formal role or support	Manual, infrequent remediation	BloodHound CE; no automation or integration
2 Managed	Small team or designated role emerging	Documented process; quarterly reviews	Pilot or partial tooling deployment
3 Defined	Dedicated Identity APM team; cross-functional coordination	Continuous remediation; integrated into changes	Enterprise-grade Identity APM tool with risk quantification
4 Quantitatively Managed	Shared accountability across key functions	Embedded in IR, detection, provisioning workflows	Tool integrated with SOAR, SIEM, ticketing systems
5 Optimizing	Embedded Identity APM ownership across architecture, IAM, security	Preventive controls during access design and policy	Policy-driven enforcement; automation blocks path creation

The Future of Identity APM as a Discipline

Identity Attack Path Management is no longer an optional side project for the red team or an occasional IT hygiene initiative. It's becoming a core discipline; on par with vulnerability management, incident response, and identity governance.

What's driving this shift is simple: no other control surface offers the same combination of visibility, impact, and preventability.

- **Identity APM sees risk the way attackers do:** through relationships, privilege chains, and abuse paths—not isolated configurations.
- **It complements your existing tools:** Identity APM strengthens PAM, IGA, EDR, and ITDR by addressing what they miss; how legitimate access is combined and misused.
- **It scales with the environment:** As identity sprawl accelerates, Identity APM helps you regain control over what privilege actually means in practice.

The most mature organizations aren't just cleaning up attack paths; they're preventing them entirely. They're using Identity APM insights during **access provisioning**, **role design**, and **system architecture**, embedding guardrails before risk takes shape.

Identity Attack Path Management is evolving into a security discipline, not just an operational function. And with the right model, teams can take a structured, measurable path to get there.

About SpecterOps

SpecterOps is a leader in identity risk management. Possessing deep knowledge of adversary tradecraft, the company enables global organizations to detect and remove critical attack paths before sophisticated attackers can take advantage of them – a practice called Attack Path Management. SpecterOps built and maintains widely used open-source security toolsets, including BloodHound, the company's foundational tool that enables attack path management in Active Directory, Entra ID and hybrid environments. BloodHound has been recommended by the U.S. Department of Homeland Security,[5] PricewaterhouseCoopers[6] and many others. BloodHound Enterprise is the company's managed SaaS for identity and security teams, allowing for attack path prioritization, remediation guidance and reporting to show improvements over time. **For more information on the benefits of an Attack Path Management practice, as well as SpecterOps and BloodHound, visit https://specterops.io**

[5] https://www.cisa.gov/news-events/directives/ed-21-02-mitigate-microsoft-exchange-premises-product-vulnerabilities
[6] https://www.pwc.co.uk/cyber-security/pdf/responding-to-growing-human-operated-ransomware-attacks-threat.pdf