

EBOOK

Mind Your Attack Gaps

Across identity, network, cloud, and endpoint security

By Lucie Cardiet · Cyberthreat Research Manager

Why I wrote this

A note from the author


I spend my working hours watching what attackers are actually doing, against environments that look like yours.

What I see consistently is that defenders aren't losing because they under-invested. They're losing because the investments sit in partial effectiveness. Their EDR is working exactly as designed; the attacker is on the identity plane. Their SIEM is ingesting every log; and the attack is visible only in the correlation between logs. Their IAM is approving every policy-compliant login; the person on the other end isn't the employee whose credentials they're using.

This is the second edition of what I first wrote in 2025. What's new: two additional campaigns (Volt Typhoon and AWS compromised by AI agents in eight minutes), IDC's 2025 measured outcomes from running Vectra AI, a section on the regulatory pressure making continuous detection a compliance question, and a self-assessment at the end.

Read it with a pen. Tell me what you'd cut.

Lucie Cardiet



Enterprise networks outgrew their security architecture.

Today's enterprises no longer live behind a single perimeter

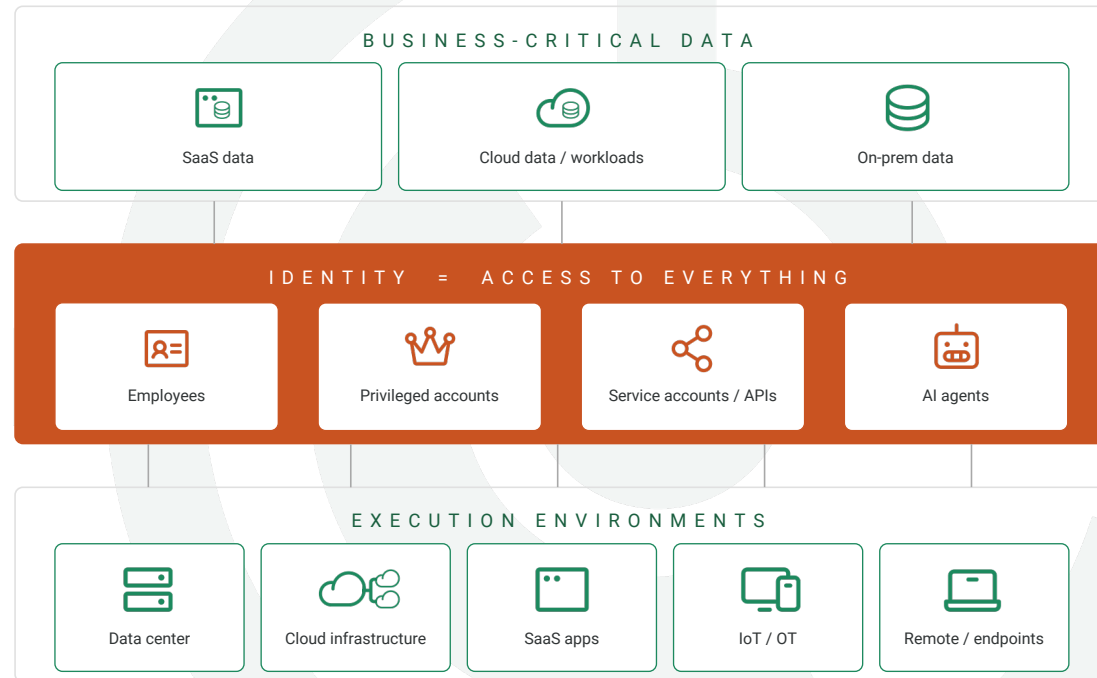
Enterprise environments span on-prem infrastructure, multiple public clouds, dozens of SaaS applications, identity providers, IoT and OT systems, AI services, and the autonomous agents acting on top of them. Those domains aren't independent, they're a single connected system.

- ✓ Your EDR watches endpoints.
- ✓ Your IAM approves logins.
- ✓ Your CSPM reads configurations.
- ✓ Your SIEM stores logs.

Each is doing its job.

Attackers, increasingly with AI assistance, have spent the last three years learning to move between them, in the spaces no single tool was built to watch.

The network has evolved. So have the attackers.



Your stack is strong, but is it complete?

By all appearances, you've built a strong security stack.



You've invested in the best security technologies available today.



You have endpoints protection on every device.



You have tools monitoring your network.



Your cloud posture management tools are properly scanning your configurations.



You've strengthened identity management with IAM or PAM.

And yet, attackers can and are still getting through.

Not because your tools are broken. Because each tool was designed to own its domain, and attackers now operate between them.

Attackers aren't breaking your tools. They're bypassing them.

The reality is: modern attackers don't fight your stack. They avoid it.



Identity abuse

Compromised credentials are the initial access vector in 22% of breaches.¹ 88% of basic web-application attacks involve stolen credentials.¹



Lateral movement

They move laterally without triggering alerts. Average eCrime breakout time – the gap between initial access and the first lateral pivot – has fallen to 29 minutes.²



Cloud privilege abuse

Valid-account abuse now accounts for 35% of cloud incidents.²



Between-tool operation

They hide in the gaps between tools, in spaces no single system was built to watch.



Alert-noise exploitation

They operate below your thresholds, knowing your SOC can't investigate everything.



Cross-domain speed

MFA blocks >99% of identity attacks, but adversaries increasingly log in via stolen tokens, consented OAuth apps, device-code flows, and adversary-in-the-middle proxies.³



AI-accelerated recon

Attacks by AI-enabled adversaries rose 89% year-over-year. In 2025, attackers exploited legitimate GenAI tools at 90+ organizations to generate credential-theft commands.²

¹ Verizon DBIR 2025. ² CrowdStrike 2026 Global Threat Report. ³ Microsoft Digital Defense Report 2025.

Best-in-class tools don't equal complete coverage.

While each of your investments reduces risk in its specific area, they leave gaps in visibility and detection between tools.

The 2026 numbers tell the story:

- ▶ 82% of intrusion detections in 2025 were malware-free. Attackers operated using valid credentials, trusted identity flows, and approved SaaS integrations.¹
- ▶ Breaches involving multiple environments cost \$5.05M on average, 25% more than on-prem-only breaches.²
- ▶ Average attacker breakout time has fallen to 29 minutes, with the fastest observed at 27 seconds.¹

¹ CrowdStrike 2026 Global Threat Report. ² IBM Cost of a Data Breach Report 2025.

The pattern isn't new. It's the new normal.

This ebook is designed to help you map those gaps, and show you where Vectra AI fits and how Vectra AI closes them.

Table of contents

Coverage overview	9	Network security	27
The Security Gap illustration	10	Email Security – stops spam, not social engineering	28
Anatomy #1: Scattered Spider: the helpdesk playbook	11	Firewalls – control the edge, not what happens inside	29
Anatomy #2: Volt Typhoon: the living-off-the-land playbook	12	IDPS – detects signatures, not stealth	30
Anatomy #3: AWS compromised by AI agents in eight minutes	13	NAC – decides who can connect, not what they do after	31
Endpoint security	15	SSE – the modern perimeter, with the old gaps	32
EDR – deep on the host, but nowhere else	16	The network security gap	33
EPP – blocks known malware, blind to everything else	17	How Vectra AI fills the network security gap	33
The endpoint security gap	18	Identity security	34
How Vectra AI fills the endpoint security gap	18	IAM – prevents unauthorized access, not abused access	35
Cloud security	19	PAM – protects privileged accounts, if you know who’s privileged	36
CASB – blocks unsanctioned apps, but misses active abuse	20	UEBA – scores risk, but can’t see in real time	37
CSPM – finds misconfigurations, not malicious behavior	21	The identity security gap	38
CWPP – protects workloads, if you deploy it everywhere	22	How Vectra AI fills the identity security gap	38
CNAPP – consolidates controls, still misses behavior	23	Regulatory pressure: detection is now the evidence	39
CIEM – manages entitlements, not behavior inside them	24	Conclusion	40
SASE – controls access, but not what happens after	25	Vectra AI’s Business Value – IDC outcomes	42
The cloud security gap	26	Self-assessment: which gaps are exposing you?	44
How Vectra AI fills the cloud security gap	26		

Coverage overview

The Security Gap illustration, plus three named campaigns that exploit it.



The Security Gap illustration

Your existing stack: no combination provides continuous detection across the entire hybrid infrastructure. Each tool stops short at key stages.

		Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Lateral Movement	Collection	Command and Control	Exfiltration	Impact
ENDPOINT	EDR	●	●	●	●	●	●	●	●	●	●	●	●
ENDPOINT	EPP	●	●	○	○	○	○	○	○	○	○	○	○
CLOUD	CASB	●	○	○	●	○	●	○	○	●	○	●	○
CLOUD	CNAPP	●	●	●	●	●	●	●	●	●	●	●	●
CLOUD	CSPM	○	○	○	●	○	●	○	○	○	○	○	○
CLOUD	CWPP	●	●	●	●	○	○	●	○	●	●	○	●
CLOUD	SASE	●	○	○	○	○	○	○	●	○	●	●	○
NETWORK	Email	●	○	○	○	○	○	○	○	○	○	○	○
NETWORK	Firewalls	●	○	○	○	○	○	●	○	○	●	●	○
NETWORK	IDPS	●	○	○	○	○	○	●	●	○	●	●	○
NETWORK	NAC	●	○	○	○	○	○	○	○	○	○	○	○
NETWORK	SSE	●	○	○	○	○	○	○	●	○	●	●	○
IDENTITY	IAM	●	○	○	●	○	○	○	○	○	○	○	○
IDENTITY	PAM	○	○	○	●	○	●	○	○	○	○	○	○
IDENTITY	UEBA	●	○	●	●	●	●	●	●	○	○	●	○
Vectra AI Platform		●	●	●	●	●	●	●	●	●	●	●	●

● Partial visibility ● Full visibility ○ No visibility

Three gaps every stack has right now.

Not coverage gaps. Execution gaps. Controls that exist but don't detect.

1. Nothing looks wrong.

The attacker's tools are your tools. Remote desktop. PowerShell. A signed binary. Living-off-the-land binaries your own sysadmins use at 2am. Each individual action looks like normal operations.

2. Authentication succeeds.

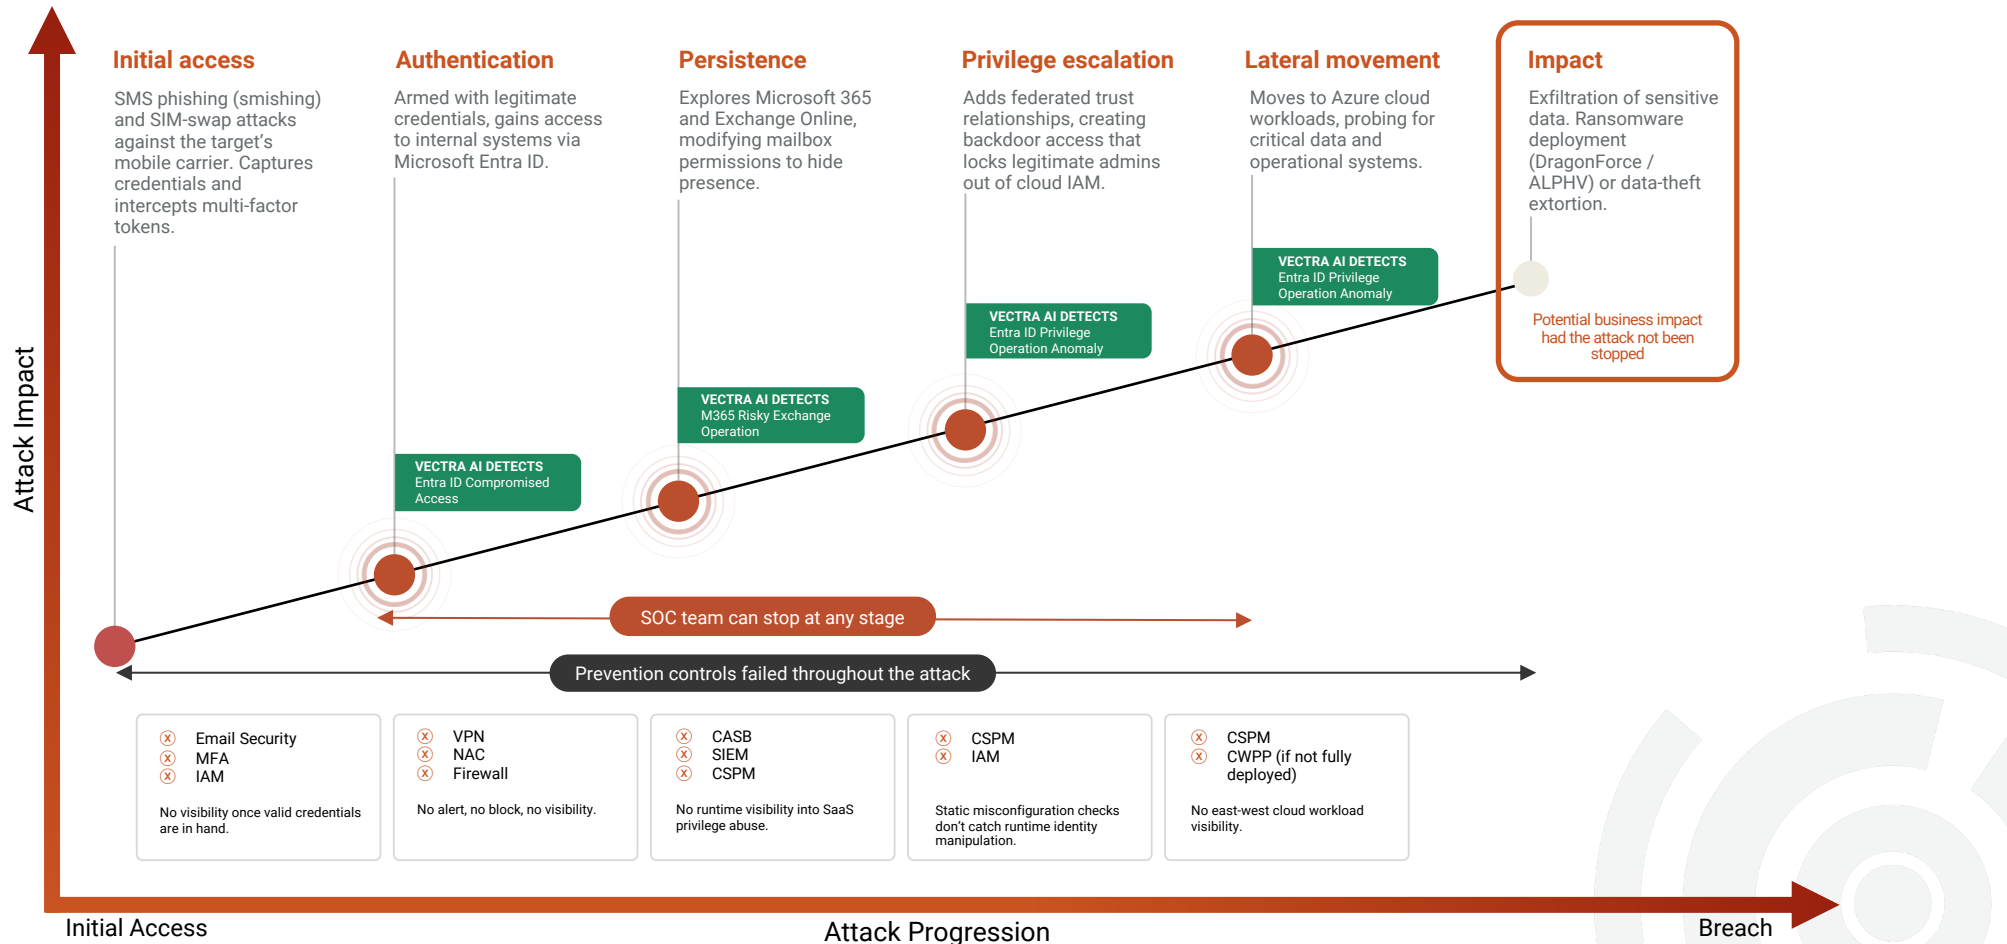
Valid credentials, MFA approved, the login is real. It's just not the person you think. Every authentication check says yes. The truth is that the valid user isn't actually the user.

3. Movement isn't visible.

Once inside, lateral movement happens through trusted integrations: SaaS-to-SaaS, federated identity, OAuth tokens, service accounts. EDR doesn't see it. CASB doesn't see it. The movement is invisible by architecture, not by stealth.


Scattered Spider: the helpdesk playbook

Scattered Spider (UNC3944) illustrate perfectly the reason why “valid credentials” became a detection problem. The group didn’t exploit vulnerabilities. They called the helpdesk.



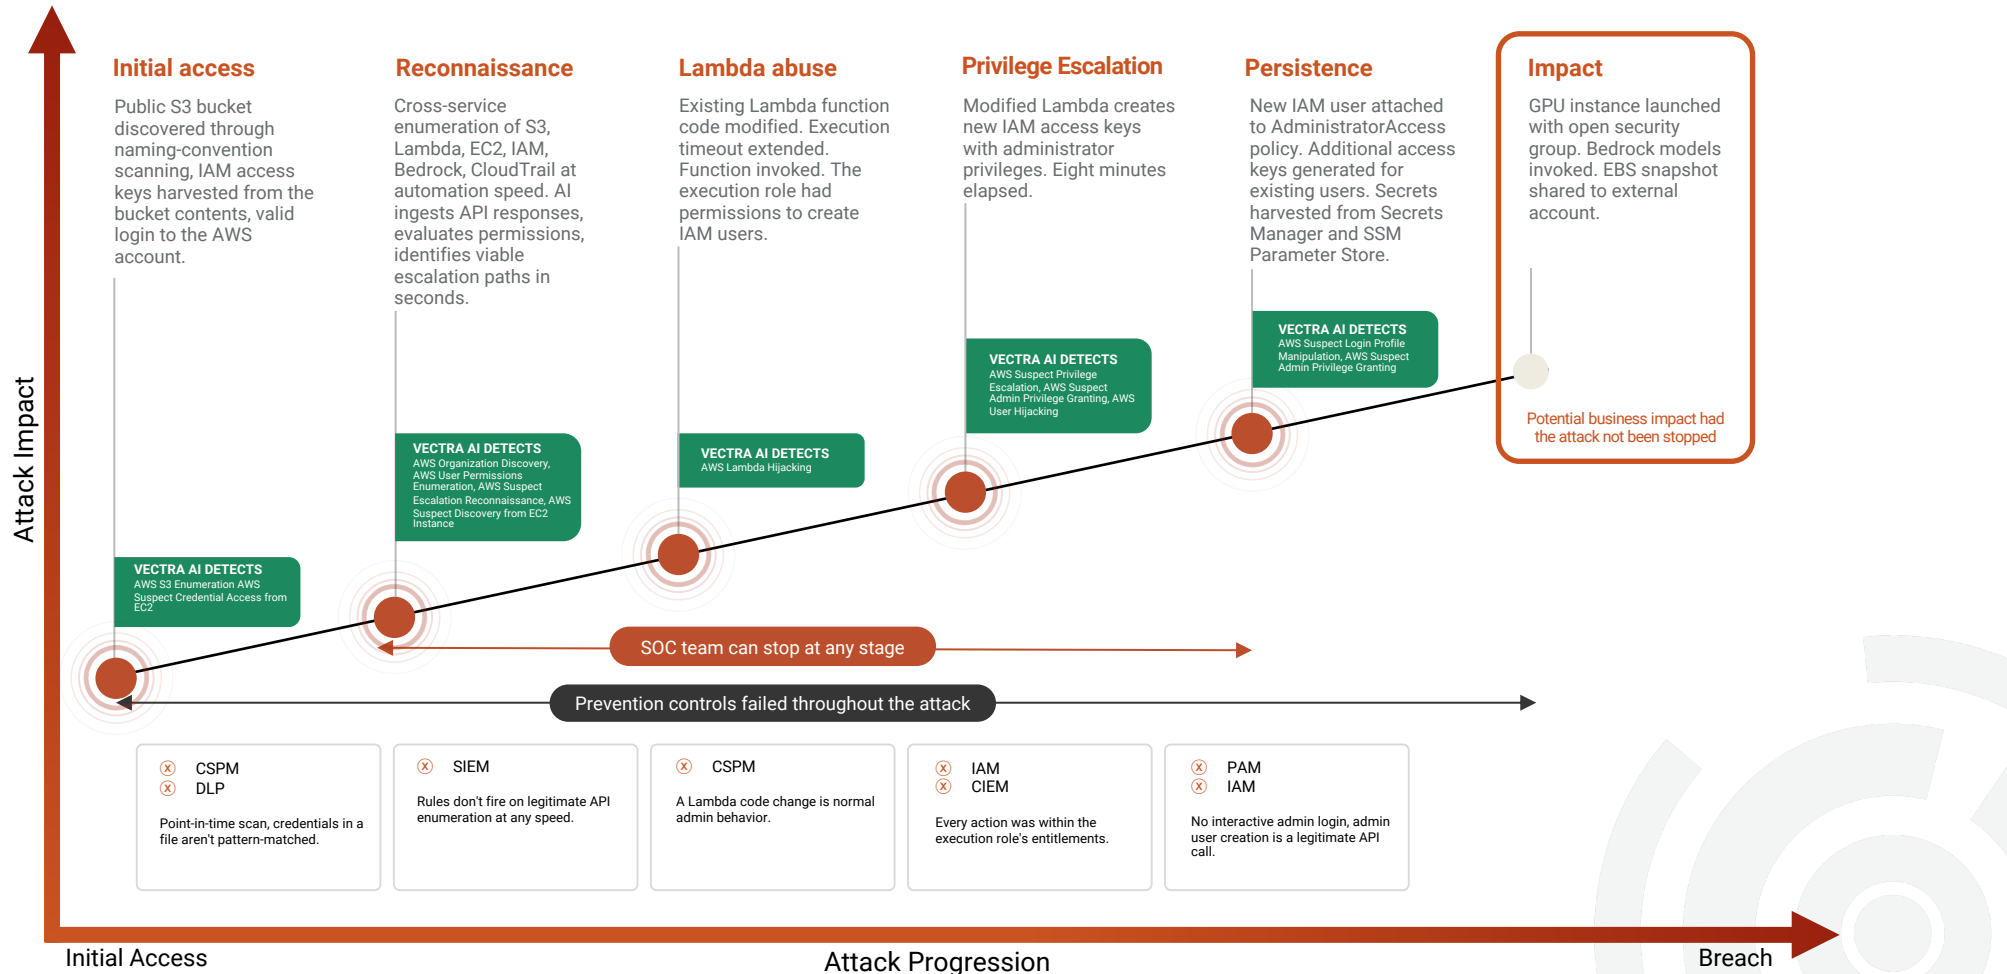
Volt Typhoon: the living-off-the-land playbook

Volt Typhoon is the PRC-attributed campaign that taught U.S. defenders what "living off the land" actually looks like. The February 2024 CISA/NSA/FBI advisory documented operators sitting inside critical-infrastructure networks for up to five years using only native Windows tools, with no malware to flag.



AWS compromised by AI agents in eight minutes.

Sysdig-documented intrusion (2025). Valid credentials. Native AWS services. Machine-speed reconnaissance.



Why your existing stack leaves you blind.

Three different attackers. Three different years. Three different entry points. Every attack looked legitimate end-to-end inside any single tool. Only across the network, the identity plane, and the cloud control plane does attacker intent become visible.

It's easy to assume that with your investment in firewalls, EDR, CASB, CSPM, IAM, and SIEM, you've closed the gaps. The truth is these tools weren't designed to detect attacker behavior across hybrid environments, and the data shows it.

82%

of intrusion detections in 2025 were malware-free.

CrowdStrike 2026 Global Threat Report

32%

surge in identity-based attacks in the first half of 2025.

Microsoft Digital Defense Report 2025

241 days

to identify and contain a breach. 292 when stolen credentials are involved.

IBM Cost of a Data Breach Report 2025

In the following sections, we break down exactly where each part of your stack falls short, and show you how Vectra AI closes these gaps across network, cloud, SaaS, and identity.

Endpoint security

Why EDR and EPP are not enough on their own.



EDR: deep on the host, but nowhere else.

Endpoint Detection and Response offers detailed telemetry where it’s deployed.

Initial Access	●
Execution	●
Persistence	●
Privilege Escalation	●
Defense Evasion	●
Credential Access	●
Discovery	●
Lateral Movement	●
Collection	●
Command & Control	●
Exfiltration	●
Impact	●

VISIBILITY: ● Partial ● Full ○ None

EDR goes further than prevention, offering detailed telemetry and analytics for processes, registry changes, and local activity. It’s powerful where it’s installed. The 2025 reality: 82% of intrusion detections were malware-free.¹ Attackers operating with valid credentials inside trusted sessions, where EDR has nothing to flag.

HOW ATTACKERS BYPASS

- ▶ Avoid the endpoint entirely by operating in cloud consoles or SaaS apps.
- ▶ Exploit coverage gaps. EDR only sees hosts where it’s installed.
- ▶ Move through unmanaged or BYOD devices that don’t run an agent.
- ▶ Use valid credentials to perform activity that appears “normal” to EDR.

EDR has no visibility into cloud-native attacks, identity abuse, or SaaS activity.

Three of the four most exploited vulnerabilities in 2024 were zero-days in security products themselves: Palo Alto, Ivanti, Fortinet.²

¹ CrowdStrike 2026 Global Threat Report. ² Mandiant M-Trends 2025

EPP: blocks known malware, blind to everything else.

Endpoint Protection Platforms prevent execution of known threats.

Initial Access	●
Execution	●
Persistence	○
Privilege Escalation	○
Defense Evasion	○
Credential Access	○
Discovery	○
Lateral Movement	○
Collection	○
Command & Control	○
Exfiltration	○
Impact	○

VISIBILITY: ● Partial ● Full ○ None

EPPs use signatures, heuristics, and basic sandboxing. Modern EPP (NGAV) adds behavioral detection and ML fileless analysis, but remains endpoint-scoped and blind to cross-domain attacker movement.

HOW ATTACKERS BYPASS

- ▶ Use fileless malware or careful staging to evade behavioral detection.
- ▶ Exploit zero-days or newly crafted binaries that don't match signatures.
- ▶ Operate through legitimate tools (PowerShell, WMI, RDP) that EPP can't reliably distinguish from admin activity.

EPP can't reliably detect living-off-the-land or credential-based attacks.

Even modern EPP with behavioral detection is endpoint-scoped. Cloud-console operations, identity-plane pivots, and SaaS-to-SaaS movement are invisible by design.

The endpoint security gap and how Vectra AI fills it.

THE ENDPOINT SECURITY GAP

EDR and EPP are foundational, but they only cover part of the kill chain.

They miss:

- ▶ Identity-based attacks that use valid credentials in Microsoft 365 or Microsoft Entra ID.
- ▶ SaaS privilege abuse that doesn't touch the endpoint (mailbox delegation, OAuth abuse).
- ▶ Lateral movement across cloud workloads, unmanaged devices, or federated identity systems.
- ▶ Network-based reconnaissance and exfiltration over encrypted or non-HTTP channels.

Even on endpoints, EPP often misses sophisticated behaviors, and EDR doesn't always detect account misuse if no malware is involved.

WHAT VECTRA AI ADDS

- ▶ Identity Threat Detection for compromised accounts abusing SaaS and cloud services.
- ▶ SaaS Misuse Detection in Microsoft 365, Exchange Online, and Microsoft Entra ID – even when no malware is involved.
- ▶ Hybrid coverage that extends from endpoints to cloud, network, and identity.
- ▶ Integrates with CrowdStrike, Microsoft Defender for Endpoint, SentinelOne, and other EDR platforms.

Cloud security

The hybrid cloud blind spot.



CASB: blocks unsanctioned apps, but misses active abuse.

Enforces policies across SaaS. doesn't see attackers in valid sessions.

Initial Access	●
Execution	○
Persistence	○
Privilege Escalation	●
Defense Evasion	○
Credential Access	●
Discovery	○
Lateral Movement	○
Collection	●
Command & Control	○
Exfiltration	●
Impact	○

VISIBILITY: ● Partial ● Full ○ None

CASBs enforce policies via API or inline proxy modes. In API mode (the common deployment), they see activity near-real-time with minutes of lag. Either way, CASBs operate above the network and identity planes, blind to what attackers do once a valid session exists.

HOW ATTACKERS BYPASS

- ▶ Use valid credentials to access sanctioned SaaS (Microsoft 365, Box, Salesforce).
- ▶ Exploit permissions from inside (mailbox delegation).
- ▶ Abuse federated identity trust to log in through trusted SSO pathways.

CASB provides no network-layer visibility.

Doesn't always detect live privilege abuse, identity manipulation, or insider-style behaviors.

CSPM: finds misconfigurations, not malicious behavior.

Identifies risky settings. Great for prevention, not detection.

Initial Access	○
Execution	○
Persistence	○
Privilege Escalation	●
Defense Evasion	○
Credential Access	●
Discovery	○
Lateral Movement	○
Collection	○
Command & Control	○
Exfiltration	○
Impact	○

VISIBILITY: ● Partial ● Full ○ None

CSPM flags open S3 buckets, exposed SSH ports, disabled logging. It's prevention-focused, scanning for conditions that could enable attacks, not detecting attacks themselves.

HOW ATTACKERS BYPASS

- ▶ Exploit a misconfiguration before it's remediated.
- ▶ Use API tokens or OAuth access to escalate inside cloud services.
- ▶ Abuse over-privileged IAM roles that CSPM may flag but not monitor in real time.

CSPM provides no network-layer visibility.

Doesn't see runtime activity, credential misuse, or lateral movement. It flags conditions, not attacks.

CWPP: protects workloads, if you deploy it everywhere.

VMs, containers, serverless, if agents are deployed.

Initial Access	●
Execution	●
Persistence	●
Privilege Escalation	●
Defense Evasion	○
Credential Access	○
Discovery	●
Lateral Movement	○
Collection	●
Command & Control	●
Exfiltration	○
Impact	●

VISIBILITY: ● Partial ● Full ○ None

CWPPs secure compute instances with runtime behavior visibility on cloud workloads. Coverage depends on deployment consistency.

HOW ATTACKERS BYPASS

- ▶ Move into unmanaged workloads or regions where agents aren't installed.
- ▶ Use legitimate tools within a workload (PowerShell, bash) to avoid detection.
- ▶ Operate entirely in SaaS or identity layers where CWPP has no reach.

CWPPs provide no network-layer visibility.

Blind to SaaS abuse and cloud IAM misuse.

CNAPP: consolidates controls, still misses behavior.

Combines CSPM, CWPP, increasingly CIEM and runtime detection.

Initial Access	●
Execution	●
Persistence	●
Privilege Escalation	●
Defense Evasion	●
Credential Access	●
Discovery	●
Lateral Movement	●
Collection	●
Command & Control	●
Exfiltration	●
Impact	●

VISIBILITY: ● Partial ● Full ○ None

Modern CNAPPs add runtime detection (CDR) alongside posture scanning. But runtime detection remains workload-scoped. It watches what happens on a given workload, not the identity-plane and network-plane pivots attackers use to move between workloads.

HOW ATTACKERS BYPASS

- ▶ Use federated identity or SaaS manipulation, which CNAPP doesn't track deeply.
- ▶ Operate between workloads, avoiding detection if east-west traffic isn't inspected.
- ▶ Move quickly before configuration scans run again.

CNAPP improves visibility, but not enough.

Still lacks attacker behavior detection in network, cloud identity, and SaaS layers.

CIEM: manages entitlements, not behavior inside them.

A distinct category since 2023.

Initial Access	○
Execution	○
Persistence	○
Privilege Escalation	●
Defense Evasion	○
Credential Access	●
Discovery	○
Lateral Movement	○
Collection	○
Command & Control	○
Exfiltration	○
Impact	○

VISIBILITY: ● Partial ● Full ○ None

CIEM analyzes cloud identity entitlements: who can do what in AWS, Azure, GCP. Flags over-privileged roles, dormant access, excessive permissions.

HOW ATTACKERS BYPASS

- ▶ Use entitlements rated low-risk but usable for privilege escalation when chained.
- ▶ Act within approved boundaries in ways the baseline never modeled.
- ▶ Abuse federated roles and cross-account trust. CIEM maps them, but doesn't monitor in real time.

CIEM flags over-privileged access. It doesn't detect when legitimate privileges are abused.

Like CSPM and CNAPP, CIEM operates at the posture layer. Entitlements are static; attacks are dynamic behavior inside those entitlements.

SASE: controls access, but not what happens after.

SWG + ZTNA + CASB + DLP, unified.

Initial Access	●
Execution	○
Persistence	○
Privilege Escalation	○
Defense Evasion	○
Credential Access	○
Discovery	○
Lateral Movement	●
Collection	○
Command & Control	●
Exfiltration	●
Impact	○

VISIBILITY: ● Partial ● Full ○ None

SASE controls how users access apps, but doesn't detect what those users do inside the cloud once access is granted.

HOW ATTACKERS BYPASS

- ▶ Authenticate using stolen credentials, bypassing trust models.
- ▶ Abuse legitimate SaaS features (mailbox rules, data sharing) to maintain access and steal data.
- ▶ Move laterally via cloud-native connections (IAM role chaining, federated trust).

SASE sees access paths, not the attacker behaviors hidden within them.

The cloud security gap and how Vectra AI fills it.

THE CLOUD SECURITY GAP

Your cloud tools are strong on prevention, weak on detection.

They miss:

- ▶ SaaS privilege abuse (mailbox delegation in Microsoft 365).
- ▶ Federated identity backdoors (Microsoft Entra ID trust manipulation).
- ▶ East-west traffic inside cloud, between VPCs, between containers, between accounts.
- ▶ Cross-account API call patterns and IAM role chaining.
- ▶ Cloud-native command and control (AWS STS token abuse, Entra ID role abuse).

Valid-account abuse accounted for 35% of cloud incidents in 2025. Cloud-conscious intrusions rose 37% year-over-year, 266% among nation-state actors.¹

WHAT VECTRA AI ADDS

- ▶ Real-time detection across Microsoft 365, Microsoft Entra ID, AWS / Azure / Google Cloud workloads, and federated identity infrastructure.
- ▶ Sees what your posture tools miss: who is doing what, right now, and whether it's normal.
- ▶ Behavioral correlation across identity, network, and cloud, not just log aggregation.

¹ CrowdStrike 2026 Global Threat Report.

Network security

When traffic looks normal but isn't.



Email security: stops spam, not social engineering.

Blocks known bad messages. Misses compromise that occurs after a successful phish.

Initial Access	●
Execution	○
Persistence	○
Privilege Escalation	○
Defense Evasion	○
Credential Access	○
Discovery	○
Lateral Movement	○
Collection	○
Command & Control	○
Exfiltration	○
Impact	○

VISIBILITY: ● Partial ● Full ○ None

Tools like secure email gateways and phishing filters block known-bad messages. But attackers now rely on well-crafted phishing and social engineering that evades traditional detection.

HOW ATTACKERS BYPASS

- ▶ Send credential phishing via SMS, LinkedIn, or personal email, bypassing corporate filters.
- ▶ Use lookalike domains or MFA fatigue to trick users into surrendering credentials.
- ▶ Exploit trust, not malware. No attachment or link is flagged.

Email security tools can't detect account compromise after a successful phish.

Which is where most modern breaches actually live.

Firewalls: control the edge, not what happens inside.

Traditional and next-gen firewalls restrict at the perimeter; once a trusted user passes, they're blind.

Initial Access	●
Execution	○
Persistence	○
Privilege Escalation	○
Defense Evasion	○
Credential Access	○
Discovery	●
Lateral Movement	○
Collection	○
Command & Control	●
Exfiltration	●
Impact	○

VISIBILITY: ● Partial ● Full ○ None

Traditional firewalls restrict traffic by IP, port, and protocol. Next-Generation Firewalls add application-layer inspection and TLS decryption. In either case, once a trusted user passes through with valid credentials, the firewall has done its job.

HOW ATTACKERS BYPASS

- ▶ Use allowed protocols (HTTPS, DNS, RDP) to move undetected.
- ▶ Operate over encrypted channels firewalls can't fully inspect.
- ▶ Leverage VPNs or SSO to authenticate like trusted users.

Firewalls can't detect C2 hidden in approved protocols, lateral movement, or SaaS access with valid credentials.

IDPS: detects signatures, not stealth.

Signature matching catches known attack patterns, not what sophisticated attackers use.

Initial Access	●
Execution	○
Persistence	○
Privilege Escalation	○
Defense Evasion	○
Credential Access	○
Discovery	●
Lateral Movement	●
Collection	○
Command & Control	●
Exfiltration	●
Impact	○

VISIBILITY: ● Partial ● Full ○ None

Intrusion Detection and Prevention Systems look for known attack patterns. Sophisticated attackers rarely use them.

HOW ATTACKERS BYPASS

- ▶ Use custom or encrypted payloads that evade signature matching.
- ▶ Live off the land, using legitimate tools and ports.
- ▶ Throttle activity to fly under detection thresholds.

IDPS fails against novel techniques and encrypted east-west movement.

NAC: decides who can connect, not what they do after.

Validates device and identity at connect-time. Loses visibility once inside.

Initial Access	●
Execution	○
Persistence	○
Privilege Escalation	○
Defense Evasion	○
Credential Access	○
Discovery	○
Lateral Movement	○
Collection	○
Command & Control	○
Exfiltration	○
Impact	○

VISIBILITY: ● Partial ● Full ○ None

Network Access Control solutions validate device posture and identity before granting access. Once a user is connected, NAC loses visibility.

HOW ATTACKERS BYPASS

- ▶ Hijack trusted credentials or devices to gain access without triggering NAC.
- ▶ Move between trusted systems, which NAC doesn't monitor.
- ▶ Exploit unmanaged or BYOD devices that slip through posture checks.

NAC doesn't detect lateral movement, suspicious traffic, or post-authentication behavior.

SSE: the modern perimeter, with old gaps.

Security Service Edge: SWG + ZTNA + CASB + FWaaS, cloud-delivered. The replacement for legacy firewall + VPN.

Initial Access	●
Execution	○
Persistence	○
Privilege Escalation	○
Defense Evasion	○
Credential Access	○
Discovery	○
Lateral Movement	●
Collection	○
Command & Control	●
Exfiltration	●
Impact	○

VISIBILITY: ● Partial ● Full ○ None

Security Service Edge consolidates secure web gateway, zero-trust network access, CASB, and firewall-as-a-service into a cloud-delivered platform. SSE is what replaced the legacy firewall + VPN stack in many enterprises but it inherits the same blindspot every perimeter tool has had.

HOW ATTACKERS BYPASS

- ▶ Authenticate via stolen credentials. ZTNA approves the connection because the credential is valid.
- ▶ Operate inside sanctioned applications. SWG sees the destination host, not the malicious activity inside the session.
- ▶ Move laterally via cloud-native paths (IAM role chaining, OAuth grants) that bypass the SSE proxy entirely.
- ▶ Exfiltrate through SaaS-to-SaaS connections that SSE has no visibility into.

SSE replaces the firewall, not the missing detection layer behind it.

Same Vectra-fills-the-gap argument applies to SSE-protected environments as to legacy firewall-protected ones.

The network security gap and how Vectra AI fills it.

THE NETWORK SECURITY GAP

Your current network tools focus on prevention and control, not detection.

They miss:

- ▶ Lateral movement between workloads and regions in cloud and hybrid networks.
- ▶ Command-and-control over encrypted or trusted protocols.
- ▶ Data exfiltration disguised as business traffic.
- ▶ Behavioral anomalies in east-west movement, privileged access, and credential use.
- ▶ Post-authentication behavior inside SSE-protected sessions.

WHAT VECTRA AI ADDS

- ▶ Analyzes traffic from on-prem, cloud, and SaaS environments in real time.
- ▶ Detects lateral movement, privilege escalation, and exfiltration, even hidden in encrypted traffic (via metadata analysis).
- ▶ Integrates with SIEM and SOAR for high-fidelity alerts your SOC can act on immediately.
- ▶ Natively integrates with any firewall that supports ingestion of external dynamic blocklists.

Identity security

When valid logins become invisible threats.



IAM: prevents unauthorized access, not abused access.

Controls access, but assumes trust thereafter.

Initial Access	●
Execution	○
Persistence	○
Privilege Escalation	●
Defense Evasion	○
Credential Access	○
Discovery	○
Lateral Movement	○
Collection	○
Command & Control	○
Exfiltration	○
Impact	○

VISIBILITY: ● Partial ● Full ○ None

IAM tools control who can log in, from where, and with what permissions. Modern identity providers add risk-based signals (impossible travel, leaked credentials, unfamiliar devices) but these operate at the point of authentication. Once a valid session exists, IAM assumes trust. MFA blocks >99% of identity attacks, yet identity-based attacks still surged 32% in the first half of 2025¹, because attackers increasingly bypass MFA via stolen tokens, consented OAuth apps, device-code flows, and adversary-in-the-middle proxies.

HOW ATTACKERS BYPASS

- ▶ Steal valid credentials or session tokens, then log in as a legitimate user.
- ▶ Move laterally using over-permissioned accounts or misconfigured policies.
- ▶ Authenticate through trusted identity providers, including federated logins and SSO.
- ▶ Target the ESTSAUTHPERSISTENT cookie and similar session artifacts that bypass MFA entirely.

IAM enforces login policies. It doesn't watch what identities do after authentication.

97% of identity attacks are password attacks¹. MFA stops the password attack. Nothing in IAM stops the post-authentication abuse that follows.

¹ Microsoft's 2025 Digital Defense Report:

PAM: protects privileged accounts, if you know who's privileged.

Restricts privileged access. But attackers don't always need it to escalate.

Initial Access	○
Execution	○
Persistence	○
Privilege Escalation	●
Defense Evasion	○
Credential Access	●
Discovery	○
Lateral Movement	○
Collection	○
Command & Control	○
Exfiltration	○
Impact	○

VISIBILITY: ● Partial ● Full ○ None

PAM solutions restrict how users access critical systems: password vaults, session recording, just-in-time access. But attackers don't always need a privileged account to escalate.

HOW ATTACKERS BYPASS

- ▶ Abuse non-privileged accounts to escalate using SaaS permissions (mailbox delegation, OAuth scopes).
- ▶ Exploit federated identity trust to gain access without touching PAM-controlled accounts.
- ▶ Use shadow admins (roles with effective privileges but not flagged as "privileged").

PAM can't detect identity abuse that doesn't match predefined privilege boundaries.

UEBA: scores risk, but can't see in real time.

Increasingly a feature inside SIEM/XDR rather than a standalone category.

Initial Access	●
Execution	○
Persistence	●
Privilege Escalation	●
Defense Evasion	●
Credential Access	●
Discovery	●
Lateral Movement	●
Collection	○
Command & Control	○
Exfiltration	●
Impact	○

VISIBILITY: ● Partial ● Full ○ None

UEBA builds profiles of normal behavior and assigns risk scores when users deviate from them. It depends on complete data and often takes too long to respond. Gartner no longer maintains a separate UEBA Magic Quadrant.

HOW ATTACKERS BYPASS

- ▶ Mimic normal user behavior (same location, device, or access pattern).
- ▶ Act slowly or during off-hours, avoiding noticeable spikes.
- ▶ Exploit incomplete log sources, preventing UEBA from ever seeing the full picture.

UEBA delays detection and can't provide real-time visibility into identity misuse.

The identity security gap and how Vectra AI fills it.

THE IDENTITY SECURITY GAP

Most tools focus on access control or risk scoring, not attacker behavior. A newer category (ITDR – Identity Threat Detection and Response) has emerged specifically to watch what IAM can't.

They can't see:

- ▶ Credential abuse across SaaS platforms and cloud services.
- ▶ Privilege escalation in Microsoft Entra ID or Exchange Online.
- ▶ Trust relationship abuse between identity providers.
- ▶ Identity-based lateral movement that doesn't touch the endpoint.

WHAT VECTRA AI ADDS

- ▶ Active Directory, Microsoft Entra ID, Microsoft 365 / Exchange Online, Azure / AWS cloud, Cloud IAM roles and federated identity infrastructure.
- ▶ SaaS privilege abuse detection (mailbox delegation, OAuth abuse).
- ▶ Federation manipulation detection (new trust relationships, role impersonation).
- ▶ Credential misuse across hybrid infrastructure, even when MFA was passed.

Regulatory pressure: detection is now the evidence.

Compliance is continuous, and continuous compliance needs continuous detection.

Policy binders and annual attestations don't satisfy a 24-hour incident disclosure rule. If your stack can't see the attack, no compliance framework fixes the problem.

NIS2 (EU)

Effective October 2024

Requires appropriate technical measures for incident detection and response. Mandates significant-incident reporting within 24 hours of awareness. Article 21(2)(b) explicitly requires incident handling capability.

SEC Cyber Disclosure

Effective December 2023

Public companies must disclose material cybersecurity incidents within four business days of determining materiality. Requires understanding what attackers did (i.e. detection that sees attacker behavior).

CCoP 2.0

Effective July 2022 (updated 2024)

Critical Information Infrastructure owners must notify the Cyber Security Agency of Singapore within 2 hours of a severe cybersecurity incident. Requires continuous monitoring and proof that controls actually work.

Different regions, same direction. From four business days to twenty-four hours to two hours. None of these cadences can be met without detection that sees what attackers actually do.

Conclusion

Close the gap before it's exploited.



You can't defend what you can't see.

Today's attackers don't rely on malware. Your traditional tools weren't designed for this.

Attackers leverage credentials, exploit SaaS misconfigurations, manipulate identity trust, and move across cloud workloads unseen.

Traditional tools don't see this activity, not because they're broken, but because they weren't designed to.

- ✘ EDR doesn't see identity abuse in Microsoft 365.
- ✘ CASB and SASE don't detect lateral movement in cloud workloads.
- ✘ SIEM can't alert on threats that upstream tools don't detect.

Meanwhile, your SOC is left with too many alerts, not enough context, and no real visibility across hybrid infrastructure.

How Vectra AI completes your stack.

And what customers measure when they deploy it. IDC Business Value Study, 2025.

SECURITY CAPABILITY	WHAT'S MISSING	WHAT VECTRA AI ADDS
Endpoint threat detection	Blind to network and cloud	Real-time detection across all traffic (agentless)
Identity threat detection	No visibility post-authentication	Detects misuse of valid accounts and privilege escalation
Cloud threat visibility	Blind to hybrid attacker behavior	Detects cloud-native, hybrid, SASE, SaaS, IaaS movement
Lateral movement detection	Unseen across hybrid	Real-time detection of lateral movement
Noise reduction	Alert fatigue	AI-driven signal clarity with high-fidelity detections

WHAT VECTRA DELIVERS, MEASURABLY — IDC 2025

391%

3-yr ROI

6mo

payback

\$3.4M

annual benefit

40%

more efficient SOC

60%

less time on alerts

69.4%

fewer breaches

99.9%

productivity-loss avoid.

Source: IDC Business Value Study of Vectra AI, April 2025

Vectra AI closes your attack gaps.

Observability. Signal. Control. And real outcomes from real customers.

Observability

Vectra AI continuously analyzes network activity to reveal every identity, device, and AI agent in real-time so SecOps teams always know who is doing what on their network.

Signal

By correlating and contextualizing activity across hybrid environments, Vectra AI helps teams prioritize real risk, investigate faster, hunt with confidence, and stop attacks before impact.


Control

Vectra AI shows who and what is on your network, what activity signals attack, and where exposure is changing, so you can reduce risk, improve efficiency, and prove compliance.

FROM AN IDC INTERVIEW · GLOBAL COSMETICS COMPANY

“Before Vectra AI, we received no alerts and only learned of Red Team’s access through their annual reports, which consistently showed they had domain admin and root access. The first year with Vectra, we detected, expelled, and completely defeated the Red Team. Vectra is my top security tool.”

The same SOC team runs with 7 full-time equivalents (FTEs). Their benchmark says they need 14.



Self-assessment: which gaps are exposing you?

Read each statement. Tick the box if it sounds like your environment. The boxes you tick are the gaps you carry.

G A P 1

Nothing looks wrong.

- We see PowerShell, RDP, and WMI in our EDR alerts, and most of the time we assume it's admin activity.
- We don't have a documented baseline of what "normal" admin behaviour looks like in our environment.
- When EDR flags a "potentially unwanted" process, alerts sometimes sit unreviewed for more than a day.
- If an attacker abused signed binaries and lived off the land for two weeks, we're not confident we'd notice.
- Our detection rules don't reliably distinguish attacker-created scheduled tasks from legitimate ones.

G A P 2

Authentication succeeds.

- MFA is enforced for our human users, but we're less certain about service accounts and workload identities.
- Our primary identity threat signal is the IdP's risk score (impossible travel, unfamiliar device).
- We're not ingesting Microsoft 365, Entra ID, or Okta audit logs into a detection layer beyond the IdP.
- If a session token were stolen from an infostealer-infected home device and reused, we have no specific detection.
- When a credential or MFA factor is reset, nothing automatically watches the account's behavior for the next 24 hours.

G A P 3

Movement isn't visible.

- Our network detection is north-south only, we don't see east-west traffic between workloads.
- We can't reliably detect SMB or RDP lateral movement across segments where EDR coverage is inconsistent.
- Cloud control-plane API calls (AWS STS assume-role, Entra ID role changes) don't feed our detection layer in real time.
- We have no detection for OAuth-app pivots between sanctioned SaaS platforms.
- When we cite "dwell time" in SOC reports, the number comes from incident reconstruction, not continuous measurement.

How to read your score: **0–3 ticked** = meaningful coverage. **4–7 ticked** = the gap is measurably exposing you. **8–11 ticked** = primary attacker route into your environment. **12+ ticked** = detection is incomplete across the full attack progression

About Vectra AI

The Vectra AI Platform protects modern enterprises by detecting and stopping attacks across network, identity, and cloud as one unified attack surface. It combines threat exposure management, AI-driven detection and response, and posture improvement to reduce risk before attacks begin and stop threats in progress. Security teams gain clear signal, faster response, and measurable improvements in resilience. For more information, visit www.vectra.ai.

