



WHITEPAPER

Mind Your Gaps in the Age of AI-Powered Attacks

The case for Attack Signal Intelligence: where modern attacks hide, how AI changes the fight, and a plain way to tell real detection from the label.

Preface

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

What I see consistently is that security teams 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 paper walks through four real intrusions, the tool categories most teams already have in place, and **the three gaps** that let attackers move through all of them anyway: nothing looks wrong, authentication succeeds, movement isn't visible. Every section ends with a check you can run today, not a feature list.

But first, one definition, because the rest of this paper depends on it.

Threat intelligence tells you what attackers have already done, somewhere else. Indicators of compromise, malware hashes, known infrastructure, named campaigns, CVE lists. It's built from other people's incidents, and it's valuable precisely because it's specific: you can block a known bad IP address or patch a known CVE. But it's necessarily retrospective. By the time something is in a threat intel feed, someone else has already been breached by it.

Attack Signal Intelligence is Vectra AI's term for something different: behavioral detection that watches what's happening on your own network and identity infrastructure right now, and turns it into prioritized, attributed signal, without needing to already know the attacker's name or their tool. It doesn't ask whether this matches a known-bad indicator. It asks whether this behavior matches how attackers operate, and it can answer that question for a technique nobody has named yet.

That distinction is the whole argument of this paper. Every gap in the following sections exists precisely because a tool built to match known-bad indicators cannot see an attacker using your own admin tools, a real credential, or a trusted integration. None of those things are indicators. They're behavior. Threat intelligence, by design, has nothing to match them against. Attack Signal Intelligence is what watches the behavior instead.

This has been a noisy year for AI claims in security marketing. If you're skeptical of vendors saying AI and meaning very little, that skepticism is earned, and I'm not asking you to set it aside. Attack Signal Intelligence isn't an autonomous agent making unilateral decisions on your network. It's explainable behavioral detection: every alert traces to specific observed activity, correlated across network, identity, and cloud, and a human on your team decides what happens next. If a vendor can't show you the behavior behind an AI claim, that's the tell to distrust it, in our product or anyone else's.

I hope you find something in here worth acting on.



Lucie Cardiet
Cyberthreat Research Manager, Vectra AI


MIND YOUR ATTACK GAPS

Enterprise networks outgrew their security architecture.

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, and 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.



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

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

- **Lateral movement:** Average eCrime breakout time has fallen to 29 minutes.²
- **Cloud privilege abuse:** Valid-account abuse now accounts for 35% of cloud incidents.²
- **Cross-domain speed:** MFA blocks over 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.²

82% of intrusion detections in 2025 were malware-free.¹ Breaches involving multiple environments cost \$5.05M on average, 25% more than on-prem-only breaches.⁴

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

¹ CrowdStrike 2026 Global Threat Report. ² Same. ³ Microsoft Digital Defense Report 2025. ⁴ IBM Cost of a Data Breach Report 2025

THE CASE FOR Attack Signal Intelligence

Attack Signal Intelligence vs. Threat Intelligence.

One tells you what's already happened somewhere else. The other tells you what's happening on your network right now.

	Threat Intelligence	Attack Signal Intelligence
Time orientation	Retrospective — campaigns and indicators already observed elsewhere	Real-time — behavior happening on your network right now
Source	Known indicators, IOCs, published reports, threat feeds	Live telemetry across identity, network, cloud, and SaaS
What it answers	"What has this group done, elsewhere, before?"	"What is happening inside my environment, right now?"
Volume	Thousands of indicators a day, most never matched to your environment	One prioritized signal, already correlated to risk
Limitation	Only as good as what's already been seen and published	Needs behavioral visibility across the whole network to work

Threat intelligence could tell you, weeks later, that Scattered Spider was targeting helpdesks.

Attack Signal Intelligence told you, in the moment, that someone had just reset an account by phone and signed in from a device that had never touched your network before, regardless of which group was behind it, or whether it had a name yet.

Knowing who attacked you matters after the fact.

Knowing what's happening right now is what stops it.

Table of contents

Coverage overview	6
The Security Gap illustration	7
Four Attack Anatomies	8
Endpoint security	11
EPP – blocks known malware, blind to everything else	12
EDR – deep on the host, but nowhere else	12
The endpoint security gap	13
Cloud security	14
CASB – blocks unsanctioned apps, but misses active abuse	15
CSPM – finds misconfigurations, not malicious behavior	15
CWPP – protects workloads, if you deploy it everywhere	16
CNAPP – consolidates controls, still misses behavior	16
CIEM – manages entitlements, not behavior inside them	17
SASE – controls access, but not what happens after	17
The cloud security gap	18
Network security	19
Email Security – stops spam, not social engineering	20
Firewalls – control the edge, not what happens inside	20
IDPS – detects signatures, not stealth	21
NAC – decides who can connect, not what they do after	21
The network security gap	22
Identity security	23
IAM – prevents unauthorized access, not abused access	24
PAM – protects privileged accounts, if you know who's privileged	24
UEBA – scores risk, but can't see in real time	25
The identity security gap	26
AI security	27
AIDR – watches AI agents and models, not the identity behind them	28
AI-SPM – finds misconfigurations, not AI misuse	28
AI Guardrails – filters prompts and outputs, not downstream action	29
Shadow AI Discovery – finds unsanctioned AI, not what it's doing	29
The AI-security gap	30
Compliance	31
How Attack Signal Intelligence maps to regulations & frameworks	32
Conclusion	33
You can't defend what you can't see	34
Vectra AI closes your attack gaps	35
Self-assessment: which gaps are exposing you?	36
Where to go next?	37

Coverage overview

The Security Gap, plus four named campaigns that exploit it.

The Security Gaps

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	●	●	●	●	●	●	●	●	●	●	●	●
	EPP	●	●	○	○	○	○	○	○	○	○	○	○
CLOUD	CASB	●	○	○	●	○	●	○	○	●	○	●	○
	CNAPP	●	●	●	●	●	●	●	●	●	●	●	●
	CSPM	○	○	○	●	○	●	○	○	○	○	○	○
	CWPP	●	●	●	●	○	○	●	○	●	●	○	●
	SASE	●	○	○	○	○	○	○	●	○	●	●	○
NETWORK	Email	●	○	○	○	○	○	○	○	○	○	○	○
	Firewalls	●	○	○	○	○	○	●	○	○	●	●	○
	IDPS	●	○	○	○	○	○	●	●	○	●	●	○
	NAC	●	○	○	○	○	○	○	○	○	○	○	○
IDENTITY	IAM	●	○	○	●	○	○	○	○	○	○	○	○
	PAM	○	○	○	●	○	●	○	○	○	○	○	○
	UEBA	●	○	●	●	●	●	●	●	○	○	●	○
AI	AIDR	○	●	○	○	○	○	●	○	●	●	●	○
	AI-SPM	○	○	○	●	○	○	●	○	○	○	○	○
	AI Guardrails / Firewalls	●	●	○	○	○	○	○	○	○	○	●	○
	Shadow AI discovery	○	○	○	○	○	○	●	○	○	○	○	○
Vectra AI Platform		●	●	●	●	●	●	●	●	●	●	●	●

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 the person you think.
3. Movement isn't visible
- Lateral movement happens through trusted integrations: SaaS-to-SaaS, federated identity, OAuth tokens, service accounts. The movement is invisible by architecture, not by stealth.

Four attack anatomies


#1 – 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.




#2 – 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 CISA documented operators sitting inside critical-infrastructure networks for up to five years using only native Windows tools, with no malware to flag.




#3 – AWS compromised by AI agents in eight minutes.

A public S3 bucket yields IAM keys, an AI agent enumerates S3, Lambda, EC2, IAM, Bedrock, and CloudTrail at machine speed, modifies a Lambda function, and creates admin-privileged IAM keys, eight minutes elapsed.



#4 – OpenAI to Hugging Face, same behavior, new operator.

In July 2026, an autonomous OpenAI model escaped a sandbox during a reduced-safeguard evaluation, chained zero-days and stolen credentials to reach Hugging Face production, and accessed limited datasets and service credentials. Hugging Face reconstructed 17,000+ events; further details remain preliminary.



Endpoint, cloud, network, and identity security: what each category misses, and the check, change, limit

Every of these four attacks 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.



EPP: blocks known malware, blind to everything else.

Endpoint Protection Platforms use signatures, heuristics, and sandboxing to stop known threats before they execute.

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) EPP won't flag.

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

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

VISIBILITY: ● Partial ● Full ○ None

EDR: deep on the host, but nowhere else.

Endpoint Detection and Response gives deep telemetry and behavioral analytics, but only on devices where it's deployed.

HOW ATTACKERS BYPASS

- ▶ Avoid the endpoint entirely by operating in cloud consoles or SaaS apps.
- ▶ Move through unmanaged or BYOD devices that don't run an agent.
- ▶ Use valid credentials to perform activity that looks legitimately normal.

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

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

VISIBILITY: ● Partial ● Full ○ None

The endpoint security gap.

82% of 2025 intrusions involved no malware.¹ EPP misses it; EDR only sees it on managed hosts. Here's what gets through:

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.

Reconnaissance and exfiltration hidden in encrypted or non-HTTP channels.

Even on managed endpoints, EPP misses sophisticated behavior, and EDR only has eyes on the host, not the identity and SaaS layers where most of this activity happens. That's the gap the check below finds, in your own environment, today.

Endpoint Security Gap Diagnostic

1 Check

Pull 30 days of EDR alerts and count how many low-severity PowerShell, WMI, or RDP flags from privileged accounts nobody's reviewed this quarter.

2 Change

Layer identity and SaaS detection on top of EDR with Attack Signal Intelligence. Extend it, don't replace it.

! Limit

Won't catch a misconfiguration that handed over access before any behavior occurred. That takes correlation across your stack, which is below.

What Vectra AI's Attack Signal Intelligence adds.

- ▶ Correlates identity, SaaS, and endpoint signal into one view, so you're not stitching together three dashboards and a guess.
- ▶ 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.

Cloud Access Security Brokers enforce policy via API or proxy mode across sanctioned SaaS applications.

HOW ATTACKERS BYPASS

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

CASB provides no network-layer visibility and doesn't detect identity misuse or insider-style behavior.

Initial Access	●
Execution	○
Persistence	○
Privilege Escalation	●
Defense Evasion	○
Credential Access	●
Discovery	○
Lateral Movement	○
Collection	●
Command & Control	○
Exfiltration	●
Impact	○
VISIBILITY: ● Partial ● Full ○ None	

CSPM: finds misconfigurations, not malicious behavior.

Cloud Security Posture Management flags risky settings, open ports, and excessive permissions for prevention, not detection.

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. It flags conditions, not attacks.

Initial Access	○
Execution	○
Persistence	○
Privilege Escalation	●
Defense Evasion	○
Credential Access	●
Discovery	○
Lateral Movement	○
Collection	○
Command & Control	○
Exfiltration	○
Impact	○
VISIBILITY: ● Partial ● Full ○ None	

CWPP: protects workloads, if you deploy it everywhere.

Secures compute instances with runtime behavior visibility on cloud workloads. Coverage depends entirely on deployment consistency.

HOW ATTACKERS BYPASS

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

CWPP provides no network-layer visibility. It flags conditions, not attacks.

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

VISIBILITY: ● Partial ● Full ● None

CNAPP: consolidates controls, still misses behavior.

Cloud-Native Application Protection Platforms combine CSPM, CWPP, and increasingly CIEM with runtime detection.

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 still lacks attacker-behavior detection across network, identity, and SaaS layers.

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

VISIBILITY: ● Partial ● Full ● None

CIEM: manages entitlements, not behavior inside them.

Cloud Infrastructure Entitlement Management analyzes who can do what across AWS, Azure, and GCP 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 doesn't monitor in real time.

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

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

VISIBILITY: ● Partial ● Full ● None

SASE / SSE: controls access, but not what happens after.

SASE unifies SWG, ZTNA, CASB, and FWaaS into one cloud-delivered edge (SSE is its security-only subset, without SD-WAN).

HOW ATTACKERS BYPASS

- ▶ Authenticate with stolen credentials; zero-trust approves the connection because the credential is valid.
- ▶ Operate inside sanctioned applications the policy already trusts.
- ▶ Move laterally via cloud-native connections (IAM role chaining, federated trust) the edge never inspects.

SASE/SSE verifies access paths, not the attacker behavior inside them.

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

VISIBILITY: ● Partial ● Full ● None

The cloud security gap.

35% of 2025 cloud incidents involved valid-account abuse. Posture tools don't catch it.

Here's what gets through:

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 & IAM abuse (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

Cloud Security Gap Diagnostic

1 Check

List every cross-account role or federated trust relationship and sort by last-reviewed date.

2 Change

Add real-time behavioral detection across Microsoft 365, Entra ID, and cloud workloads, correlated with identity in one alert queue.

! Limit

Won't catch a genuinely novel cloud-service misuse pattern with zero prior baseline. That takes the correlation below.

What Vectra AI's Attack Signal Intelligence adds.

- ▶ Correlates identity, network, and cloud activity into one view, not three consoles and a guess.
- ▶ 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.

Network security

When traffic looks normal but isn't.



Email security: stops spam, not social engineering.

Secure email gateways and phishing filters block bad messages, not the compromise that follows a successful click.

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, so no attachment or link is ever flagged.

Email security tools can't detect account compromise after a successful phish, which is where most modern breaches actually live.

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

VISIBILITY: ● Partial ● Full ● None

Firewalls: control the edge, not what happens inside.

Traditional and next-gen firewalls restrict traffic by IP, port, and protocol, then stop inspecting once a user is trusted.

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.

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

VISIBILITY: ● Partial ● Full ● None

IDPS: detects signatures, not stealth.

Intrusion Detection and Prevention Systems look for known attack patterns, not sophisticated attacker behavior.

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.

Initial Access	●
Execution	○
Persistence	○
Privilege Escalation	○
Defense Evasion	○
Credential Access	○
Discovery	●
Lateral Movement	●
Collection	○
Command & Control	●
Exfiltration	●
Impact	○
VISIBILITY: ● Partial ● Full ○ None	

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

Network Access Control validates device and identity posture before granting access, then loses visibility once connected.

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.

Initial Access	●
Execution	○
Persistence	○
Privilege Escalation	○
Defense Evasion	○
Credential Access	○
Discovery	○
Lateral Movement	○
Collection	○
Command & Control	○
Exfiltration	○
Impact	○
VISIBILITY: ● Partial ● Full ○ None	

The network security gap.

It takes 241 days on average to identify and contain a breach.² Prevention and control aren't detection. Here's what gets through:

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.

Even SIEM and SOAR are only as good as the alerts they receive, and today's attackers make sure those alerts never come.

Network Security Gap Diagnostic

1 Check

Ask whoever owns your firewall or SSE logs what your top ten internal data-volume flows looked like last week, without a report request.

2 Change

Add behavioral traffic analysis across on-prem, cloud, and SaaS, feeding SIEM or SOAR as high-fidelity signal.

! Limit

Won't stop exfiltration over a channel you've allowed and never inspect. That needs the correlation below.

What Vectra AI's Attack Signal Intelligence adds.

- ▶ Correlates on-prem, cloud, and SaaS traffic into one view, not another feed for SIEM to drown in.
- ▶ Detects lateral movement, privilege escalation, and exfiltration, even hidden in encrypted traffic (via metadata analysis).
- ▶ Integrates natively with SIEM, SOAR, and any firewall that ingests external dynamic blocklists, so your SOC acts on signal, not noise.

Identity security

When valid logins become invisible threats.



IAM: prevents unauthorized access, not abused access.

Identity and Access Management governs authentication and permissions, checking risk signals at login, then assuming trust.

HOW ATTACKERS BYPASS

- ▶ Steal valid credentials or session tokens, then log in as an over-permissioned, legitimate user.
- ▶ Authenticate through trusted identity providers, including federated logins and SSO.
- ▶ Target persistent SSO/MFA cookies and similar artifacts that bypass MFA.

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

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

VISIBILITY: ● Partial ● Full ● None

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

Privileged Access Management restricts privileged access through vaults, session recording, and just-in-time grants.

HOW ATTACKERS BYPASS

- ▶ Abuse non-privileged accounts to escalate using SaaS permissions (mailbox delegation, OAuth scopes).
- ▶ Use shadow admins, roles with effective privileges that were never flagged as privileged.
- ▶ Move through credential trust that PAM-controlled accounts never extended to.

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

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

VISIBILITY: ● Partial ● Full ● None

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

User and Entity Behavior Analytics builds behavior profiles and assigns risk scores, increasingly a feature inside SIEM/XDR.

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, so UEBA never sees the full picture.

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

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

VISIBILITY: ● Partial ● Full ○ None

The identity security gap.

Identity-based attacks surged 32% in H1 2025.³ Access control alone can't see the misuse.

Here's what gets through:

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.

ITDR tools are built to watch exactly this, but coverage still varies widely by vendor. That's the gap the check below finds, in your own environment, today.

Identity Security Gap Diagnostic

1 Check

Pick your five most-privileged service accounts and check if anything flagged their behavior right after their last credential reset.

2 Change

Add behavioral detection across Active Directory, Entra ID, Microsoft 365, and cloud IAM, covering abuse even when MFA passed.

! Limit

Won't stop an insider with legitimate privilege who never deviates from normal. Everything else needs the correlation below.

What Vectra AI's Attack Signal Intelligence adds.

- ▶ Correlates identity signal across Active Directory, Entra ID, Microsoft 365, and cloud IAM, instead of one more isolated alert.
- ▶ Detects SaaS privilege abuse (mailbox delegation, OAuth abuse) and federation manipulation (new trust relationships, role impersonation).
- ▶ Flags credential misuse across hybrid infrastructure, even when MFA was passed.

AI security

Why watching the model isn't watching the account.



AIDR: watches AI agents and models, not the identity behind them.

AI Detection and Response tools monitor model behavior, prompt misuse, and agent actions inside AI pipelines and MLOps tooling.

HOW ATTACKERS BYPASS

- ▶ Operate through the agent's own cloud identity and permissions, which AIDR doesn't correlate to cloud IAM behavior.
- ▶ Chain legitimate-looking agent actions (enumerate, read, write) that resemble normal automation.
- ▶ Move laterally through connected services once the agent's credentials are issued, outside the pipeline AIDR watches.

AIDR watches the model. It doesn't watch what the model's cloud identity does next.

Initial Access	None
Execution	Partial
Persistence	None
Privilege Escalation	None
Defense Evasion	None
Credential Access	None
Discovery	Partial
Lateral Movement	None
Collection	Partial
Command & Control	Partial
Exfiltration	Partial
Impact	None

VISIBILITY: Partial Full None

AI-SPM: finds AI misconfigurations, not AI misuse.

AI Security Posture Management discovers AI services, builds an AI bill of materials, and flags over-permissioned or misconfigured AI resources.

HOW ATTACKERS BYPASS

- ▶ Exploit a flagged-but-unremediated misconfiguration before the next posture scan runs.
- ▶ Use an AI service's own over-permissioned role to act, not just exist.
- ▶ Operate inside a shadow AI instance the scan hasn't discovered yet.

AI-SPM provides no runtime visibility. It flags conditions, not attacks.

Initial Access	Partial
Execution	Partial
Persistence	None
Privilege Escalation	Partial
Defense Evasion	None
Credential Access	None
Discovery	Partial
Lateral Movement	None
Collection	None
Command & Control	None
Exfiltration	None
Impact	None

VISIBILITY: Partial Full None

AI Guardrails: filters prompts and outputs, not downstream action.

Guardrail and firewall tools screen model inputs and outputs in real time for prompt injection, data leakage, and policy violations.

HOW ATTACKERS BYPASS

- ▶ Craft prompts that pass the filter but still trigger an unsafe downstream tool call.
- ▶ Exploit tool-call chains the guardrail inspects individually, not in sequence.
- ▶ Operate through MCP servers or plugins outside the guarded request path.

Guardrails inspect the conversation. They don't inspect what the agent does with its own credentials afterward.

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

VISIBILITY: ● Partial ● Full ○ None

Shadow AI Discovery: finds unsanctioned AI, not what it's doing.

Shadow AI discovery platforms scan devices, browsers, and SaaS tenants for unapproved AI tools, agents, and API tokens.

HOW ATTACKERS BYPASS

- ▶ Use AI capabilities already embedded inside sanctioned SaaS tools, which don't register as "new."
- ▶ Launch on newly released agent platforms faster than discovery coverage expands.
- ▶ Operate within an approved tool's OAuth scope, which discovery treats as sanctioned by default.

Discovery tells you an agent exists. It doesn't tell you what the agent did once it had access.

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

VISIBILITY: ● Partial ● Full ○ None

The AI security gap.

AI-enabled adversary attacks rose 89% year-over-year.⁴ Model-layer tools don't watch the identity behind them.

Here's what gets through:

Agentic cloud action (AI agents act as cloud identities, with cloud permissions attackers can abuse)

Attacks that pass guardrails but still trigger unsafe tool calls downstream.

Shadow AI usage (unapproved copilots, agents, and API tokens outside any Inventory)

Cross-tool blind spots (no single AI tool correlates model, identity, and cloud behavior together.)

Even where AIDR, AI-SPM, guardrails, and discovery tools are deployed, each watches one layer. None of them follow the agent's credentials into the cloud, network, or identity systems it touches next.

AI Security Gap Diagnostic

1 Check

Pull your AI-SPM or shadow-AI inventory and count how many discovered agents have never had their cloud actions reviewed.

2 Change

Correlate agent and model activity with the cloud, network, and identity telemetry you already collect, don't treat AI as a separate silo.

! Limit

Won't catch an agent misusing a brand-new permission structure no one's modeled yet. That takes the correlation below.

What Vectra AI's Attack Signal Intelligence adds.

- ▶ Sees agent and model activity as identity and workload behavior, not a category requiring a separate stack.
- ▶ Detects the lateral movement and privilege abuse that happens after a guardrail, AI-SPM scan, or discovery tool has already done its job.
- ▶ One correlated signal across human and agent actors, same detection logic either way.

Compliance

Detection is now the evidence.

How Attack Signal Intelligence maps to regulations & frameworks

Here’s how Attack Signal Intelligence maps to the frameworks that matter most across the world.

What it requires			Govern	Identify	Protect	Detect	Respond	Recover
EMEA REGULATIONS	DORA – Incident Reporting	4-hour classification; initial, intermediate, and final incident reports to the competent authority.	●	●	●	●	●	○
	DORA – Risk Management	Continuous ICT risk management framework across the financial entity’s network.	●	●	●	●	○	○
	NIS2 Directive	24-hour early warning, 72-hour detailed notification, one-month final report.	○	●	○	●	●	○
	ECB SSM	Supervisory expectations for cyber resilience and incident reporting at significant institutions.	●	●	●	●	●	○
	GDPR	72-hour breach notification to the supervisory authority; data protection and accountability obligations across the EU.	●	●	●	●	●	○
	PSD2 & PSD 3	Secure electronic payments, strong customer authentication, and fraud-prevention obligations for payment providers.	○	●	●	●	●	○
	EBA ICT & Security Guidelines	ICT governance and operational resilience expectations for EU banks – the supervisory baseline DORA later formalized into law.	●	●	●	●	●	○
US REGULATIONS	SEC Item 1.05 (8-K)	Material cybersecurity incident disclosure within 4 business days of a materiality determination.	○	●	○	●	●	○
	SEC Item 106 (10-K)	Annual disclosure of cybersecurity risk management, strategy, and governance.	●	●	●	●	○	○
	CIRCIA	72-hour reporting of substantial cyber incidents; 24-hour reporting of ransom payments, to CISA.	○	●	○	●	●	○
	CMMC v2	Cybersecurity maturity certification for DoD contractors handling Controlled Unclassified Information (CUI).	●	●	●	●	●	○
	PCI DSS 4.0	Logging, monitoring, and access control requirements for any organization handling payment card data.	○	●	●	●	●	○
	NYSDFS (23 NYCRR 500)	Cybersecurity program, risk assessments, MFA, and 72-hour incident notification for NY regulated financial institutions.	●	●	●	●	●	○
	FFIEC	Risk assessments, access controls, and continuous monitoring expectations for US financial institutions.	●	●	●	●	●	○
	CDM Phase 3 DEFEND	Continuous diagnostics, asset visibility, and incident response for US federal civilian agency networks.	○	●	●	●	●	○
	DFARS / NIST 800-171	Cybersecurity controls and incident reporting for DoD contractors handling CUI; works alongside CMMC.	○	●	●	●	●	○
APJ REGULATIONS	Australia SOCI Act	Critical Infrastructure Risk Management Program; 72-hour ransom-payment reporting.	●	●	●	●	●	○
	Singapore Cybersecurity Act	Expanded CII reporting, including interconnected non-CII and third party systems.	●	●	○	●	○	○
	Japan Active Cyber Defense Law	New reporting, cooperation, and operational-security duties, phased through 2026–2027.	●	●	○	●	○	○
FRAMEWORKS	NIST Cybersecurity Framework (CSF 2.0)	Voluntary structure organizing practice into Govern, Identify, Protect, Detect, Respond, and Recover – the same six functions this table uses.	●	●	●	●	●	○
	MITRE ATT&CK	A knowledge base of adversary tactics and techniques used to benchmark detection coverage, not a compliance mandate.	○	●	○	●	●	○
	CIS Critical Security Controls (v8)	Prioritized safeguards for preventing, detecting, and responding to the most common attacks; a voluntary best practice baseline.	○	●	●	●	●	○

● Direct ● Supporting ○ Not covered

Conclusion

Which gap can you least see?

A large, light green curved shape that starts from the bottom right and curves upwards and to the left, partially overlapping the dark blue background.

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

Today's attackers don't rely on malware. They lean on 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.

What customers measure when they deploy Attack Signal Intelligence

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 was also named a Leader in the 2025 and 2026 Gartner® Magic Quadrant™ for Network Detection and Response, recognized for both completeness of vision and ability to execute.

Vectra AI closes your attack gaps.

Observe behavior, signal risk, take control before impact. Real outcomes from real customers.

Observability

Attack Signal Intelligence 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, Attack Signal Intelligence helps teams prioritize real risk, investigate faster, hunt with confidence, and stop attacks before impact.


Control

Attack Signal Intelligence 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 gap can you least see?

Nine questions, three per gap. Read each one honestly, then jump to the note for whichever gap you flagged most, that's the one to go check first.

Gap 1 – Nothing looks wrong

- 1 Can you say, right now, what your baseline of “normal” admin PowerShell, RDP, or WMI activity looks like, or would you have to guess?
- 2 If an attacker lived off the land using signed binaries for two weeks, would anyone notice before the two weeks were up?
- 3 When EDR flags something “potentially unwanted,” how long does it typically sit before a human looks at it?

Gap 2 – Authentication succeeds

- 1 Is MFA enforced for your service accounts and workload identities the same way it is for human users, or is that still a gap?
- 2 If a session token were stolen from an infected personal device and reused, would you have a specific detection for that, or only the identity provider's general risk score?
- 3 When a credential or MFA factor gets reset, does anything automatically watch that account's behavior for the next 24 hours?

Gap 3 – Movement isn't visible

- 1 Does your network detection cover east-west traffic between workloads, or only north-south at the perimeter?
- 2 Do cloud control-plane calls — an AWS assumed role, an Entra ID role change — feed into a detection layer in real time?
- 3 When you cite “dwell time” in a SOC report, is that number from continuous measurement, or reconstructed after the fact from an incident?


Where to go next

If this ebook was useful, here's where to go next:




eBook Telling Real Security AI From the Label

How AI actually works, how attackers are using it, and a plain three-question test for any vendor's AI claim, including ours.




Attack Lab Ep. 1 Nothing Looks Wrong

Living-off-the-land tradecraft: Volt Typhoon, Salt Typhoon, Flax Typhoon, and BRICKSTORM, side by side.



Attack Lab Ep. 2 Authentication Succeeds

Identity-driven access abuse: Scattered Spider, OAuth supply-chain attacks, and the Snowflake intrusions.



Attack Lab Ep. 3 Movement Isn't Visible

Cross-plane lateral movement, including the autonomous AI agent that compromised Hugging Face.



Attack Lab Ep. 4 Defending at AI Speed

When an intrusion that used to take days compresses into minutes, your response has to compress too.



Newsletter Mind Your Attack Gaps

A new case like the ones in this ebook, every two weeks, with a specific check attached.



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.