

WHITE PAPER

Research Report on Identity Threat Detection and Response (ITDR) Requirements

VECTRA®

Research Report on Identity Threat Detection and Response (ITDR) Requirements

Section 1 Identity is the new perimeter for both attackers and defenders

Section 2 Industry's perspective on ITDR Requirements

Section 3 Vectra AI's perspective on ITDR Requirements

Section 4 Recommendations and Next Steps


Section 1

Identity is the new perimeter for both attackers and defenders

In today's world, identity is at the heart of every enterprise and hybrid attack. Attackers are constantly exploiting identities to launch and spread their attacks, whether it's to get a foothold or to move around within a network to access sensitive data and deploy ransomware.

With services like phishing-as-a-service and ransomware-as-a-service, attackers have made it easy to replicate identity-based attacks on a large scale. Even with millions of dollars spent on security tools, 90% of organizations have experienced an identity attack.

This is huge number that continues to grow. But why?



90% of organizations have experienced an identity attack.

Identity is no longer locked away

Identities used to be accessible only if an attacker was already on the network. An attacker could do a lot with this access, but there was a sense that firewalls, EDR, and policies were the front lines for identities. This is not the case anymore; identities are outside the perimeter, and they are everywhere. Identities with unified cloud and network access are the norm today, accessible outside the traditional network environment, making them the new front line for defenders.

The hybrid environment is complex

As organizations continue to migrate to the Cloud, their environments span on-premises infrastructure, cloud services, and remote workspaces, creating a complex fabric of interconnected systems. This gives attackers multiple entry points to begin the attack. A single compromised entry point can lead to significant breaches as attackers are known to pivot between on-premises and cloud environments.

Blind spots: Machine and Service Identities

Often overlooked, machine identities (such as APIs, bots, and service accounts) pose unique challenges. Unlike human users, they cannot authenticate via MFA. Yet, they have access to critical resources. According to Silverfort, 31% of all users are service accounts with high access privileges and low visibility. Additionally, on average, 109 new shadow admins are introduced by a single AD misconfiguration, enabling attackers to reset a true admin's password.[†]

Siloed detection in hybrid attacks

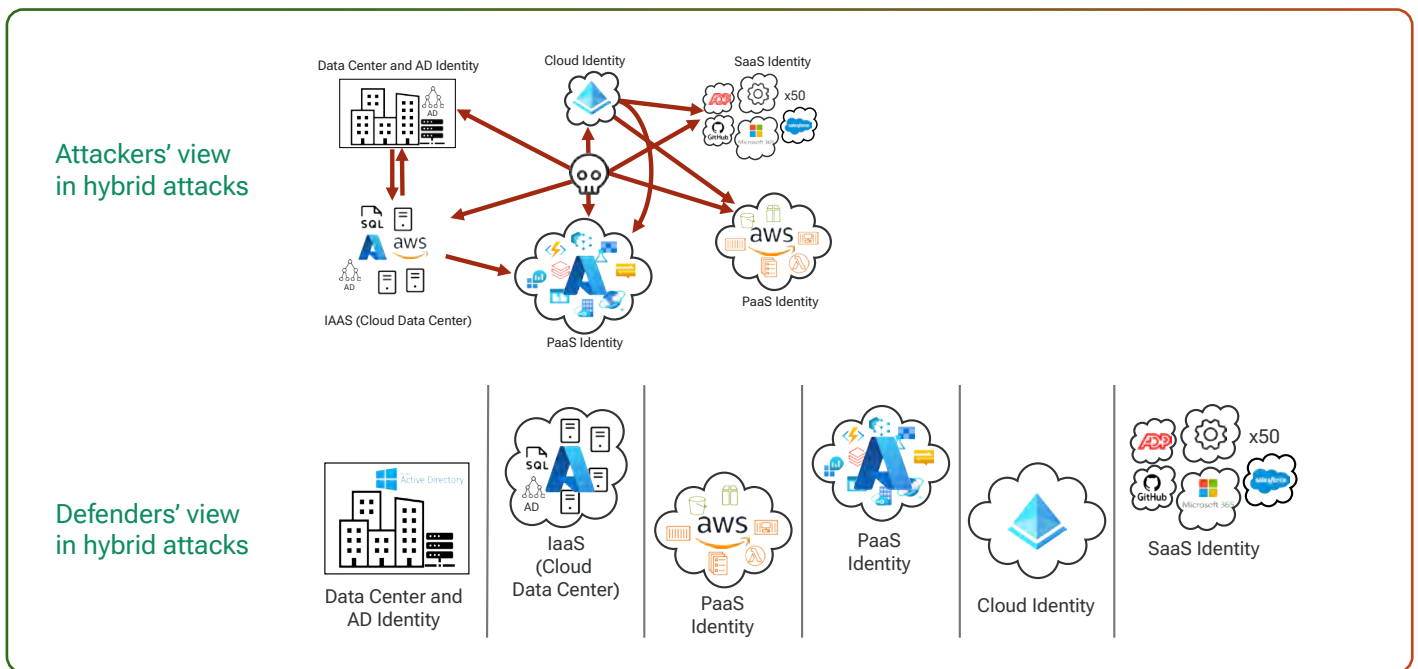
While attacker groups constantly pivot between attack surfaces in the network and cloud, defenders often rely on siloed tools for different attack surfaces. This creates challenges to identify attackers in their environment.

Expanding identity attack surface

The number of identities security teams need to protect is huge and only getting bigger – making the identity attack surface a huge target for attackers. Every user (customers, employees, partners and vendors), device, and service account in the cloud and network represent a potential attack vector. According to IDSA, 98% of organization saw an increase in identities and 62% don't have visibility into the humans and machines accessing sensitive data and assets.*

Attackers need two things to be successful – an identity and a network

Attacker groups like Scattered Spider or ALPHV/ Black Cat are leveraging many different tactics to exploit legitimate access credentials to infiltrate companies and persist in organizations' environments. With the rapid increase in enterprise identities – users and machines – coupled with the lack of visibility, attackers have more ways to get on the network and progress their attacks.



Prevention Isn't Foolproof: MFA and EDR limitations

While preventive measures like Multi-Factor Authentication (MFA) and Endpoint Detection and Response (EDR) play crucial roles, they can be bypassed. Attackers can bypass MFA through social engineering or compromised devices.

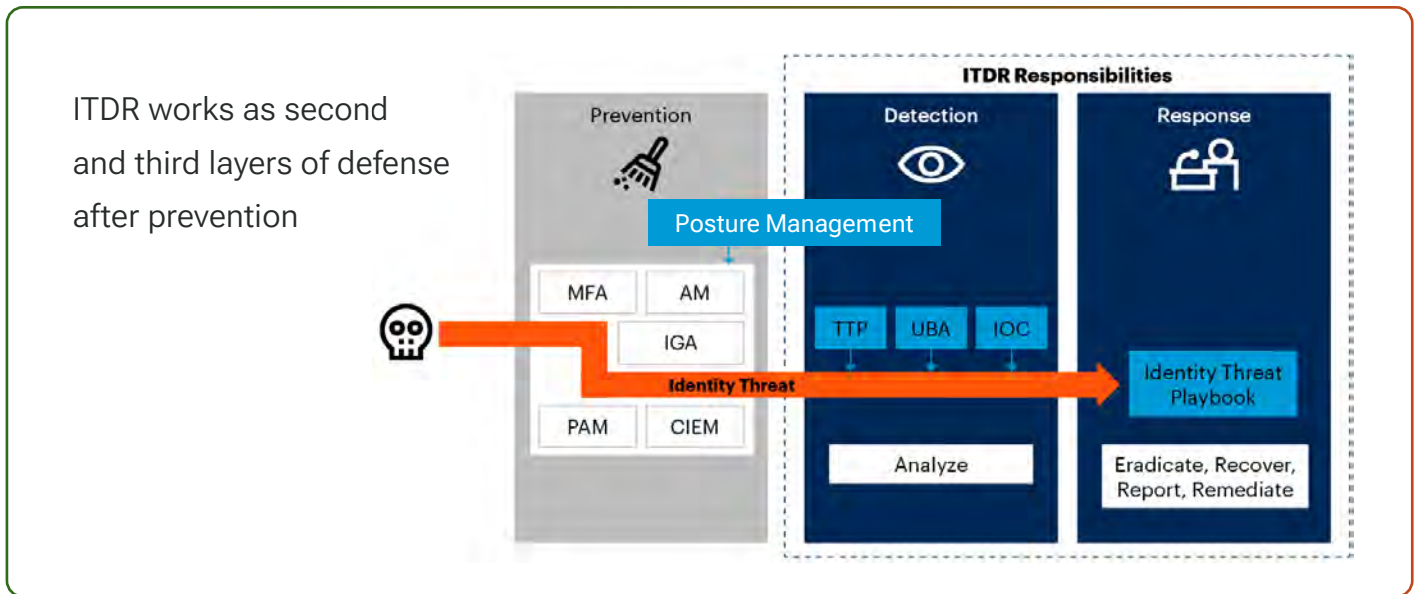
Similarly, EDR solutions may miss subtle signs of identity compromise. Organizations need to augment prevention with robust detection and response capabilities.

Gen AI attack surface – increasing identity exposure exponentially

GenAI tools, like Copilot for Microsoft 365, aim to increase efficiency and productivity for employees. However, they also create new attack surfaces because the large language models (LLMs) powering these tools have access to proprietary corporate data. While GenAI tools enhance the speed at which employees and organizations operate, they

also provide adversaries with the same advantage. This enables attackers to increase the speed and scale at which they exploit identities to spread their attacks. Organizations require detection and monitoring capabilities to prevent and stop attackers from abusing identities through GenAI tools and accessing sensitive data and information.

The introduction of Identity Threat Detection and Response (ITDR)



Vectra AI has observed that attackers have proven ways to bypass prevention and Identity posture tools. According to Gartner®, “ITDR works as second and third layers of defense after prevention”¹

In this report, we delve into the industry’s perspective on ITDR requirements. We study the common features in ITDR solutions available in the market and evaluate the gaps they might have in defending against identity attackers in hybrid environments.

Section 2

Industry's perspective on ITDR requirement

According to Gartner[®], “Good ITDR must detect changes in the security posture (part of hygiene), and detect ongoing attacks.”² Vectra AI has observed that there are two categories of ITDR solutions in the market: posture-focused solutions aimed at preventing compromise, and post-compromise-focused solutions aimed at detecting attackers already in the environment. Most ITDR solutions focus on enhancing identity posture to prevent identity attacks, with some offering partial post-compromise capabilities. Below is a list of commonly observed capabilities for these two categories of ITDR solutions.

“Good ITDR must detect changes in the security posture (part of hygiene), and detect ongoing attacks.”²

GARTNER

Posture focused/ Pre-compromise	Post-compromise
<ul style="list-style-type: none"> • Hybrid Identity coverage (Human and non-human identities across on-premises and cloud) • Integration with existing toolsets (e.g. SIEM, SOAR, XDR, etc.) and Identity provider (IDP) 	
<ul style="list-style-type: none"> • Posture assessment <ul style="list-style-type: none"> • Tracks the security posture of the identity attack surface and make recommendations to enhance identity hygiene • Assess identities vulnerability • Proactive protection <ul style="list-style-type: none"> • Proactive capabilities to reduce threat and remediate vulnerability before an attack • Compliance management <ul style="list-style-type: none"> • Assess Identity threats against defined compliance and governance standards, e.g. MITRE ATT&CK • Risk Scoring and Prioritization <ul style="list-style-type: none"> • Prioritize and remediate the riskiest identity vulnerabilities • Prioritize anomalies in user authentication and access activities • Detect changes in security posture <ul style="list-style-type: none"> • Monitor and analyze alternations to security configuration and settings • Monitor limited initial access and lateral movement identity attack techniques <ul style="list-style-type: none"> • Focus on meeting checkbox requirement • Automated/ Customized incident response <ul style="list-style-type: none"> • Apply controls to isolate attacks and stop their spread • Central dashboard <ul style="list-style-type: none"> • Provide an overview of risk analysis of identity posture and hygiene • Central alerting <ul style="list-style-type: none"> • Alert on anomalies in user activities regarding authentication and access attempt • Advanced reporting <ul style="list-style-type: none"> • Report such as benchmarking, risk analysis, audit logging and risk scoring of vulnerable identities 	<ul style="list-style-type: none"> • Real time threat detection and response post compromise when attackers bypass prevention <ul style="list-style-type: none"> • Detects attacker behaviors • Deception techniques that trap attackers • Correlated signal <ul style="list-style-type: none"> • Correlate information from different attack surface across different types of identities and the network to give early warning signals and visibility • Risk Scoring and Prioritization <ul style="list-style-type: none"> • Prioritize compromised identity factoring in entity importance • AI minimize noise • Monitor identity attack techniques e.g. MITRE ATT&CK <ul style="list-style-type: none"> • Cover MITRE ATT&CK techniques throughout the attacker kill chain (including initial access, living-off-the-land, and exfiltration techniques) • Automated/ Customized incident response <ul style="list-style-type: none"> • Apply controls to isolate attacks and stop their spread • Incident response analysis <ul style="list-style-type: none"> • Provide investigation logs and metadata for security teams • Central dashboard <ul style="list-style-type: none"> • Provide an overview of compromised identities and attacker techniques observed in customers' environment • Central alerting <ul style="list-style-type: none"> • Alert customers about threats in their environment based on customized threshold • Advanced reporting <ul style="list-style-type: none"> • Report such as benchmarking, risk analysis, audit logging and risk scoring of identity attacks

The question is: Are these enough to stop identity attacks? Here's our observation:

Which is more important: attacker that might get in or attacker already in?

According to [Vectra AI research](#), 71% of SOC analysts think they are already compromised, and they just don't know it yet. Posture and hygiene are important and fundamental to identity-first security. However, it is an ongoing work-in-progress with new users, devices, systems, and workloads. Not to mention misconfigurations that arise from automations or system changes due to M&A activities. Posture, processes, and hygiene take time to get right. It's

a never-ending activity, with always more gaps to close and more configurations to change. Despite the best efforts to close every gap, attackers only need one opening to progress in an environment. As attackers continue to accelerate their speed in attacks, security teams should prioritize investment in post-compromise threat detection so that they can stop attackers who have already infiltrated their environment as early as possible, before damage occurs.

No one checks all the boxes

There are very few solutions that focus on both pre-compromise and post-compromise in the market. It is important for security teams to evaluate what they have in their toolkits, maximize the capabilities of their existing identity technology first, and adopt a solution that fills the gap of missing capabilities while integrating well with the existing technology stack.

Breadth of detection - Correlation is key

Different solutions may specialize in the coverage of various attack surfaces and types of identities, but it is essential for security teams to ensure that they have visibility into both human (customers, employees, vendors, partners) and non-human identities (workloads, devices) across data centers, networks, and cloud environments. Additionally, since attackers pivot within organizations' hybrid environments, it is crucial to correlate information across different attack surfaces to detect early warning signals of attacker behaviors. Siloed detections create a vulnerable point for attackers to hide within your environment.

Depth of detection - Security testing is essential

While different solutions may claim that they can detect attack behaviors and compromised identities, it is important for security teams to compare the depth of detections along the attack path, whether the detection is based on anomalies or attacker behaviors. One way to validate this is through security testing. Instead of treating security testing as a checkbox exercise, security teams are recommended to leverage online tools to conduct attack simulations and battle-test their security environments' resilience against real attacker behaviors.



Does the AI help prioritize the risk while reducing the noise?

While AI is being used in various technologies, security teams should prioritize solutions that help identify true threats and minimize noise. Signal clarity is the ultimate path to enable SOC teams to focus on what is urgent and real, allowing them to get ahead of attackers and stop attacks.

Section 3

Vectra AI's perspective on important criteria of ITDR

With the above observations, below are some important criteria that security teams should consider when selecting the right ITDR tool. We categorize these criteria into Coverage, Clarity and Control.

Coverage that reduces exposure

- Post compromise threat detection
- Detection of real attacker techniques, not anomaly
- Detect changes in security posture to enhance identity hygiene
- Detection of GenAI tools abuse

Clarity that removes latency

- AI prioritization that cut through the noise
- Integrated and correlated visibility in human, non-human, cloud and network identities

Control that maximizes talent

- Stop identity attacks at speed
- Cut investigation times
- Investigation context and content
- Comprehensive response capabilities
- Integration with existing toolkits



ITDR Solution

Coverage that reduces exposure

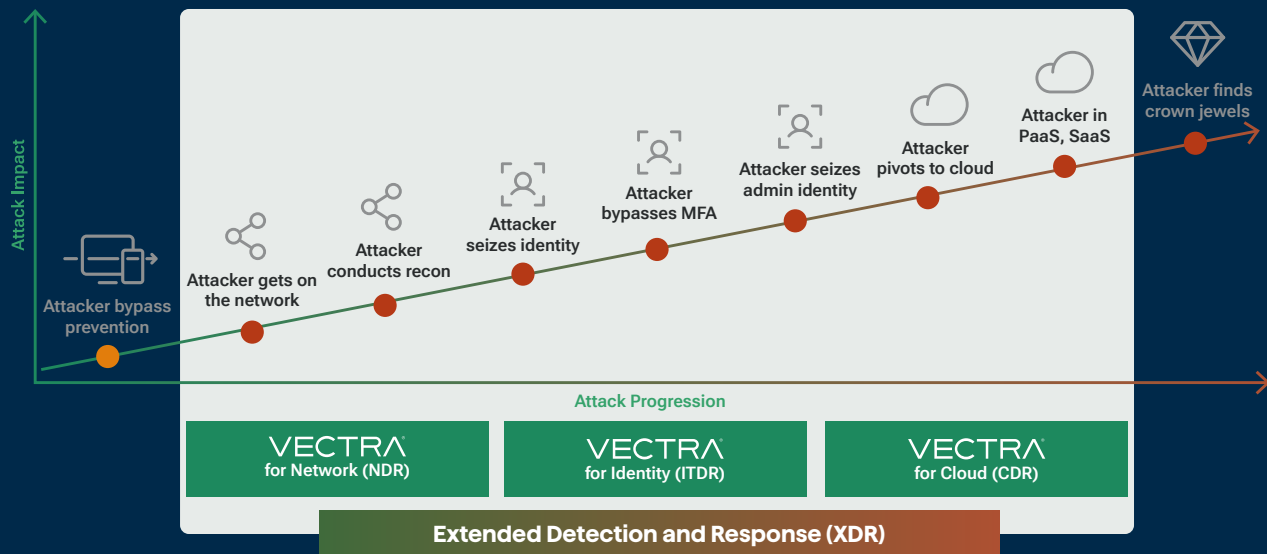
Post compromise threat detection is key

Preventative controls like multi-factor authentication (MFA) are crucial for protecting identities, however, they can be bypassed. Once an attacker infiltrates your environment, the only defense is monitoring their actions. SOC teams must detect compromised users based on identity behavior within the network and correlate actions across domains.

As attackers increasingly abuse identities to log in rather than hack in, identifying those misusing identities is vital for SOC teams.

The Vectra AI Platform covers over 90% of MITRE ATT&CK techniques, including post-compromise methods, helping security teams reduce attack exposure.

Vectra AI coverage reduces attack exposure post compromise



Vectra AI covers over 90% of MITRE ATT&CK techniques

Access	Persist	Command & Control	Escalate & Evade	Recon & Discover	Lateral Movement	Exfiltration & Disruption
New Host	MFA Disabled	Hidden HTTPS Tunnel	New Host Role	Kerberoasting (x2)	Privilege Access Anomaly (x6)	Smash and Grab
Suspected Compromise Access	Trusted IP Change	Hidden DNS Tunnel	Log Disabling Attempt	Internal Darknet Scan	Suspicious Remote Exec	Ransomware File Activity
Brute-Force Attempt/Success	Admin Account Creation	Hidden HTTP Tunnel	Disabling Security Tools	Port Scan	Suspicious Remote Desktop	Data Gathering
Disabled Account	Account Manipulation	Multi-homed Fronted Tunnel	Suspicious Mailbox Rule	Port Sweep	Suspicious Admin	Data Smuggler
TOR Activity	Redundant Access	Suspicious Relay	Log Disability Attempt	SMB Account Scan	Shell Knocker	Hidden DNS Tunnel Exfil
Unusual Scripting Engine	Logging Disabled	Suspect Domain Activity	Suspect Privilege Escalation	Kerberos Account Scan	Automated Replication	Hidden HTTP/S Tunnel Exfil
Suspicious OAuth App	User Hijacking	Malware Update	Suspect Privilege Manipulate	Kerberos Brute-Sweep	Brute-Force	Botnet Abuse Behaviors
Suspicious Sign-On	ECS Hijacking	Peer-to-Peer	Suspect Console Pivot	File Share Enumeration	SMB Brute-Force	Crypto Mining
Suspicious Sign-On with MFA Fail	Suspect Login Profile Manipulation	Suspicious HTTP	Suspect Cred Access EC2	Suspicious LDAP Query	Kerberos Brute Force	External Teams Access
Suspicious Teams App	Security Tools Disabled	Stealth HTTP Post	Suspect Cred Access SSM	RDP Recon	Suspect Cred Access SSM	Ransomware SharePoint Activity
Suspicious Credential Usage	SSM Hijacking	TOR Activity	Suspect Cred Access ECS	RPC Recon	Internal Stage Loader	Suspicious SharePoint Download
Root Credential Usage		Novel External Port	Suspect Cred Access Lambda	RPC Targeted Recon	Suspicious Active Directory	Suspicious SharePoint Sharing
TOR Activity		Threat Intel Match		Unusual eDiscovery Search	Novel Admin Protocol	Exfil Before Termination
		Vectra Threat Intel Match		Unusual Compliance Search	Novel Admin Share Access	Suspicious Mailbox Forwarding
				Suspect eDiscovery Activity	Risky Exchange Op	eDiscovery Exfil
				User Permission Enumeration	Internal Spear Phishing	Power Automate Activity (x3)
				EC2 Enumeration	File Poisoning	Ransomware S3 Activity
				S3 Enumeration	Mailbox Manipulation	Suspect Public S3 Change
				Suspect Escalation Recon	DLL Hijacking	Suspect Public EBS Change
				Organization Discovery	Privilege Operation Anomaly	Suspect Public EC2 Change
						Suspect Public RDS Change
						Suspect External Access Grant

- Data Center Network & Identity, IaaS, IoT/OT
- Identity: Active Directory, Azure AD
- PaaS: AWS
- SaaS: Microsoft 365

Detection of real attacker techniques, not anomaly

To counter hybrid attacks by groups like Scattered Spider or Midnight Blizzard (APT29) and proactively address unknown and ever-evolving attacker techniques (including zero days and novel attacks), security teams require behavior-based detections that comprehend underlying attacker behaviors. Simple rule-based user and entity behavior analytics (UEBA) fail to identify attacks and overwhelm analysts with alerts.

Vectra AI has pioneered AI-driven attack signal detection within data over the past decade. Our security research team delved into attacker mindset and analyzed revealing data. As

a result, we've developed 150+ AI models for various attack types. These models drive real-time analytics, achieving high recall and precision. For instance, Vectra's patented graph-based AI algorithm, Privilege Access Analytics, continuously monitors critical accounts, services, and hosts, uniquely identifying privilege abuse. The AI algorithm calculates observed privilege by considering the historical interactions between tracked entities, rather than relying solely on the privilege defined by an IT admin. This approach enables more accurate detection of privilege escalations.

Detect changes in security posture to enhance identity hygiene

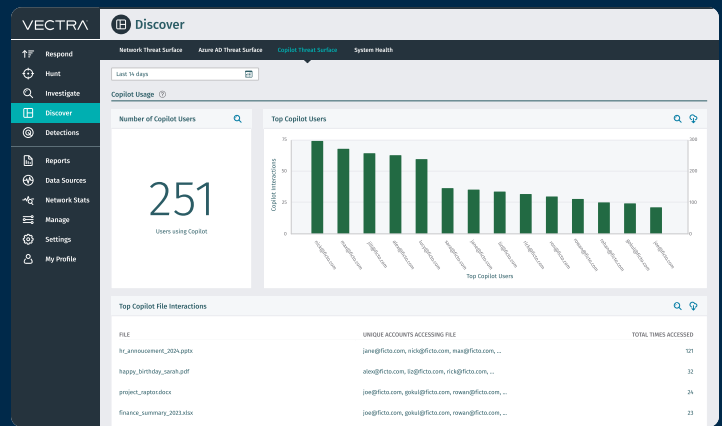
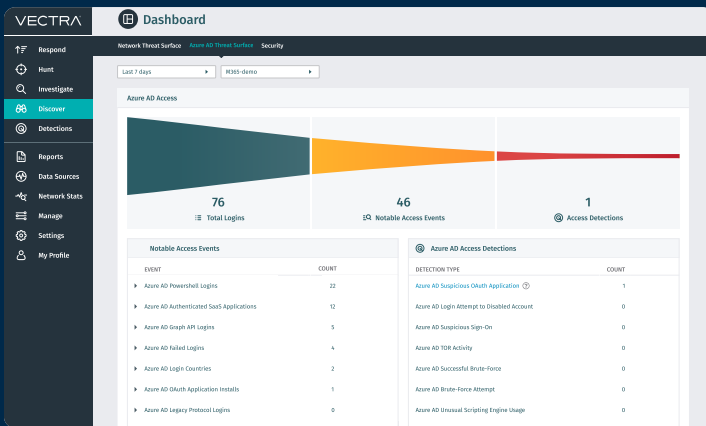
In order to reduce the attack surface related to identity, companies should proactively and continuously monitor their security posture to identify any weaknesses and vulnerabilities. It is important to have a tool that generates key actionable insights to improve identity hygiene.

The Vectra AI Discover Dashboard highlights Identity hygiene issues like account logins without two-factor authentication, use of legacy sign-in protocols, overly permissive access to tools like PowerShell and graph API,

and weak location-based access controls. This enables proactive monitoring and identity hygiene enhancement.

To pinpoint identity hygiene issues, Vectra AI's investigation capabilities provide access to underlying metadata and visibility into normal user behavior. It identifies areas where static configuration controls are not followed. For instance, it detects accounts not using MFA for logins, logins from locations that should be blocked by conditional access policies, and access abuse related to admin privileges.

Vectra AI Discover Dashboard highlights Identity hygiene issues



Detection of GenAI tools abuse

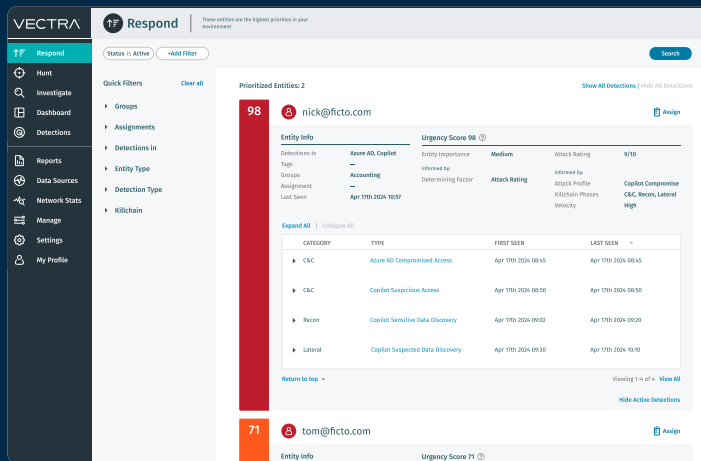
As enterprise begin to adopt GenAI like Copilot for Microsoft 365 (Over 40% of Vectra AI's identity customers have begun the adoption of Copilot for M365), it is essential for organizations to have visibility into identity's usage of GenAI tools so that they can prevent and stop attackers from abusing these tools to access and exfiltrate sensitive information and data.

Vectra AI's new AI-driven detections for GenAI attacks empower SOC teams to fight AI with AI, enabling them to operate at the same speed and scale as attackers. The Vectra AI Platform detects attackers that compromise an identity and abuse Copilot for Microsoft 365 to execute their attacks and gain access to sensitive applications and

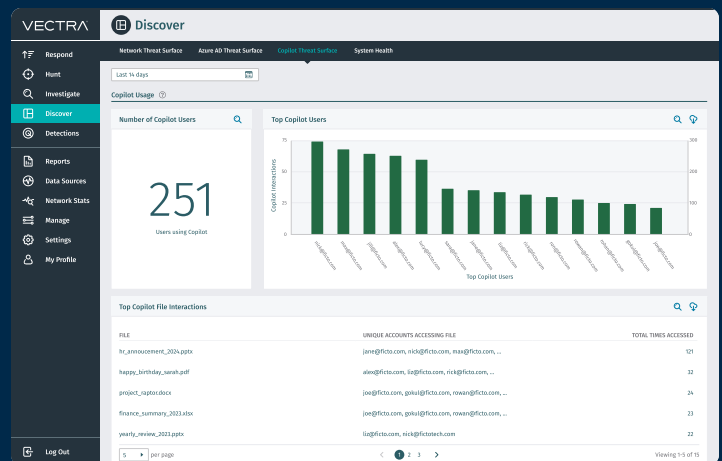
data in a matter of minutes. To achieve this, the Vectra AI Platform delivers:

- GenAI detection coverage for Microsoft Copilot abuse including suspicious access, data discovery and jailbreak attack techniques.
- Correlation and attribution of GenAI detections with identities across Microsoft Entra ID, Microsoft 365, AWS, and Active Directory.
- Gain visibility into users' Copilot for M365 usage on Discover Dashboard, including top users and data access for auditing and security posture improvement.

Vectra AI Detects Copilot for Microsoft 365 abuse



Vectra AI Discover Dashboard highlights Copilot for Microsoft 365 usage



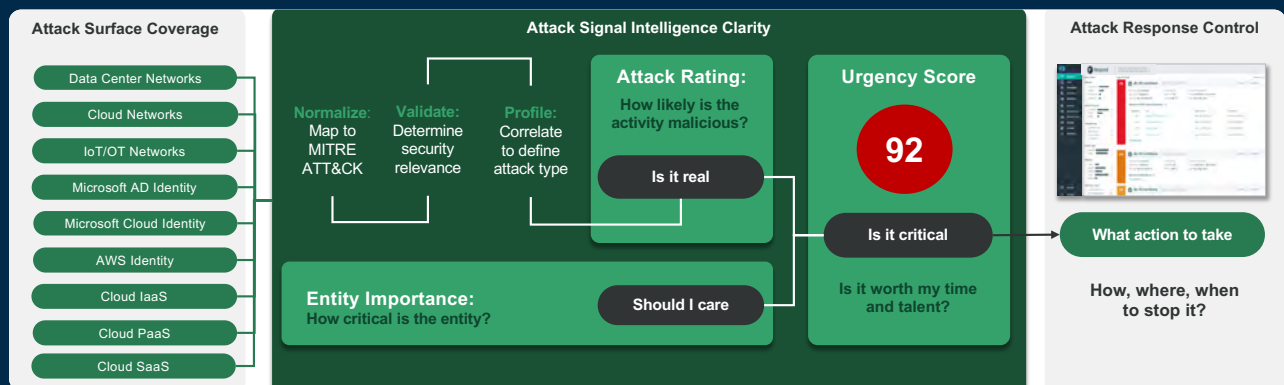
Clarity that removes latency

Cut through the noise, once and for all

Even if you have all the alerts and intelligence that you need, do you have the manpower and time to consume all data, sift through the false positives to get the relevant information? An effective and efficient ITDR should have a threat prioritization engine that cut through the noise and only show you relevant attack vectors that the team should care about.

Vectra AI helps customer reduce 100% noise compared to traditional UEBA solutions. Together with the flexibility to create customized prioritization rules, this provides security teams the clarity on real attacker behaviors so that they can free up man hours to do other productive tasks.

Vectra AI's AI-driven Attack Signal Intelligence delivers accurate and integrated attack signal to SOC analysts



Integrated visibility in human, non-human, cloud and network identities

To effectively detect hybrid attacks, security teams require visibility into both human and non-human identities in their cloud and on-premises environments. This is especially crucial for non-human identities such as machine accounts, service accounts, instance credentials, service principles, or application credentials, for which MFA cannot be enabled due to legacy systems and automation requirements. Additionally, an effective ITDR solution should correlate all relevant information to provide early warning signals to defenders.

Vectra AI's identity detection covers network Active Directory, Microsoft Entra ID, Microsoft 365, and AWS IAM identities. The Vectra AI Platform correlates identity activities with broader network and cloud activities, providing integrated visibility so defenders can identify attackers in their environment and stop privilege escalation and lateral movement.

Control that maximizes talent value

Comprehensive response capabilities

An effective ITDR solution should provide analysts peace of mind, even in the middle of the night or on weekends, by offering automatic response capabilities to stop attackers swiftly and at scale, around the clock, and customized response options to ensure flexibility in security operations.

Vectra AI's comprehensive response capabilities offer native, integrated, and managed responses. Native response allows security teams to automatically lock down accounts and revoke sessions, disabling compromised

identities or associated endpoint systems on the Vectra AI Platform. Customized response options allow security teams to investigate identities while minimizing operational disruption. Integrated response enables security teams to isolate threats that reach a certain threshold on the Vectra AI platform via SOARs and EDRs. Additionally, managed response supports security teams with experienced MXDR analysts who offer remote response and remediation using Microsoft Defender, SentinelOne, and CrowdStrike EDRs.

Investigation context and content are critical

Investigation content and context provide the necessary insights for security teams to investigate potential threats and understand the root cause of incidents.

Vectra ITDR provides query-less access to investigation data across network Active Directory, Microsoft Azure

AD, Microsoft 365, and AWS IAM identities, which is not commonly provided by other tools. This offers an easy way for SOC analysts to conduct proactive investigations and perform effective post-incident analysis.

Integration with existing toolkits

The ITDR solution should fit into your existing environment and integrate with your toolkits. This will ease adoption and help make the solution a core part of the security technology stack.

Vectra AI's Attack Signal Intelligence provides integrated, accurate hybrid attack signals to your pane of glass,

eliminating the need to change SOC processes and workflows. Integrations with IAM, SIEM, SOAR, or EDR/XDR technologies reduce cost and complexity, minimize latency, shorten learning curves, and accelerate time to value.



Section 4

Recommendation

While different organizations are at varying levels of maturity in identity security, Vectra AI recommends that Security Leaders, Architects, Engineers, and Analysts consider the following:

Prioritize post-compromise focused ITDR first for SOC

Security teams are recommended to develop processes to detect attackers who are already in your environment. While posture is important, a balanced approach is recommended to reduce the identity attack surface. As tools consolidation continues to be a priority, security teams should leverage native tools' best practices on hygiene before acquiring a tool that enhances identity posture. Additionally, complementing this approach with a post-compromised focused ITDR is essential.

Evaluate vendors by outcome-based criteria

When evaluating vendors, security teams should consider outcome-based criteria:

- Coverage that reduces exposure
- Clarity that removes latency
- Control that maximizes talent and technology

By adopting this approach, security teams can realistically prevent and stop identity breaches efficiently and effectively without overburdening their existing manpower and resources.

Involve multiple stakeholders in purchase decision

When ITDR is being purchased, the SOC should be at the table when making purchase decisions, together with the GRC and IAM teams. This will ensure a balanced perspective is considered and maximize the benefit of the tool for different teams.

Next steps

Conduct exposure gap analysis

IAM and GRC teams would benefit from conducting a gap analysis to identify gaps in achieving balanced pre-compromise and post-compromise identity coverage in their tools.

[Calculate the risk of your organization's identity exposure](#)

Conduct red team exercises that simulate attack paths used by the latest hybrid attackers

SOC teams should perform these exercises to enhance their understanding of resilience against attacks by hybrid groups like Scattered Spider and Midnight Blizzard. If budget and resources are limited, leveraging free online tools is recommended.

[Sign up for a guided security test](#)

Assess current tools

Security architects and detection engineers are recommended to assess the organization's identity security toolsets and select a tool that can easily integrate with their existing toolsets while addressing identity security gaps.

* [IDSA The State of Identity and Security](#)

† [Silverfort Identity Underground Report](#)

1 Gartner Conference, Technical Insights: Mitigate Identity Threats with Identity Threat Detection and Response, Paul Rabinovich, 4 June 2024

2 Gartner Conference, Improve Enterprise Resilience with Identity Hygiene, Security Posture Management and ITDR, Rebecca Archambault, 3 June, 2024

GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved.

About Vectra AI

Vectra AI is the leader in hybrid attack detection, investigation and response. The Vectra AI Platform delivers integrated signal across public cloud, SaaS, identity, and data center networks in a single platform. Vectra AI's patented Attack Signal Intelligence empowers security teams to rapidly detect, prioritize, investigate and stop the most advanced hybrid cyber-attacks. With 35 patents in AI-driven detection and the most vendor references in MITRE D3FEND, organizations worldwide rely on the Vectra AI Platform and MXDR services to move at the speed and scale of hybrid attackers. For more information, visit www.vectra.ai.

For more information please contact us:

Email: info@vectra.ai | vectra.ai

© 2024 Vectra AI, Inc. All rights reserved. Vectra, the Vectra AI logo, and Security that thinks are registered trademarks and the Vectra Threat Labs, Threat Certainty Index and Attack Signal Intelligence are trademarks of Vectra AI. Other brand, product and service names are trademarks, registered trademarks or service marks of their respective holders. Version: 072924