

How Vectra AI delivers Zero Trust visibility and security capabilities

The initial point of contact during a cyberattack is rarely the intended target. Attackers usually gain access to networks from a less secure workstation or IoT asset and work their way from there by gaining access to higher privileged hosts and accounts.

Zero trust

This is why the concept of zero trust has grown significantly in the last couple of years, since cloud applications and a mobile workforce have redefined the security perimeter. Modern corporate resources and services now often bypass on-premises, perimeter-based security models that rely on network firewalls and VPNs, and have become obsolete. A zero-trust architecture fundamentally considers all entities in a network to be hostile and does not allow any access to resources until both the account and host have been individually authenticated and authorized to use that specific resource.

The gaps with access-only approaches

However, this approach to zero trust commonly seen in Privileged Access Management and Identity Access Management solutions still relies on single point in time-security gating decisions that use a predefined list of privileged identities. There are several issues with this approach being the sole implementation.

One issue involves simple configuration errors. This is especially common in cloud environments due to the differentiated skillset required to manage the complexity of constantly changing cloud resources as opposed to traditional on-prem counterparts.

Another issue is that once granted, access can easily be manipulated by attackers who use methods like credential abuse and privilege escalation. Both of these methods are especially hard for security practitioners to detect. They seldom have any visibility into the credentials being used on the network versus credentials assigned by Identity Providers (IdPs).

Key Highlights:

- Monitoring what interactions are actually occurring on the network – rather than looking at application logs or assigned permissions – exposes the immutable truth of what goes on inside a specific environment.
- The Vectra AI Platform continuously monitors the behaviors of users, hosts and services, and applies supervised and unsupervised AI models to score these behaviors for threat, certainty and prioritization of risk.
- Vectra currently leverages more than 98% of the MITRE ATT&CK framework, and Vectra AI is constantly including more with every software release.

Continuous visibility and assessment of privilege

Closing this gap requires extending the preliminary method of authentication and authorization by continuously monitoring what accounts and identities are being used to access the network and cloud. This is why Visibility and Analytics has been the driving component of Forrester's Zero Trust Ecosystem.

According to Gartner, "security and risk management leaders need to embrace a strategic approach where security is adaptive, everywhere, all the time. Gartner calls this strategic approach "continuous adaptive risk and trust assessment," or CARTA and

"with a CARTA strategic approach, we must architect for digital business environments where risk and trust are dynamic and need to be assessed continuously after the initial assessment is performed."

With monitoring, it is possible to observe if behaviors deviate from expectations in a risky way, and surface this to security practitioners to determine if access to the capabilities should be adapted or removed entirely.

Monitoring what interactions are actually occurring exposes the immutable truth of what goes on inside a specific environment.

The Vectra AI approach

The Vectra AI Platform uses AI to efficiently find and prioritize hidden attacks in real-time inside your cloud services like Microsoft Office 365, Azure AD, cloud, data center, IoT, and enterprise networks before attackers cause irreparable harm to the organization. The platform allows security teams to prevent attacks earlier in the kill chain, ensuring that applications essential to business continuity are available and accessible for the entire extended workforce.

As a key component of a Zero Trust Framework, Vectra AI will help deliver visibility and analytics on three guiding principles:

1. Verify explicitly. Always authenticate and authorize based on all available data points, including user identity, location, device health, service or workload, data classification, and anomalies.
2. Use least privileged access. Limit user access with Just-in-Time and Just-Enough (JIT/JEA), risk-based adaptive policies, and data protection to protect both data based adaptive policies, and data protection to protect both data and productivity.
3. Assume breach. Minimize blast radius for breaches and prevent lateral movement by segmenting access by network, user,

devices, and application awareness. Verify all sessions are encrypted end to end. Use analytics to get visibility, drive threat detection, and improve defenses.

In order to achieve this visibility, the following entities need to be fully observed:

- The hosts that are accessing workloads
- The servers or services that contain the workloads
- The user or service accounts being leveraged.

The Vectra AI Platform continuously monitors the behaviors of accounts, hosts and services, and applies supervised and unsupervised AI models to score these behaviors for threat, certainty and prioritization of risk.

As a result, Vectra AI delivers a continuous real-time assessment of privilege. This empowers security teams with the right information to anticipate what assets will be targeted by attackers, and to rapidly take action against the malicious use of privilege across cloud and hybrid environments.

Flexible AI models

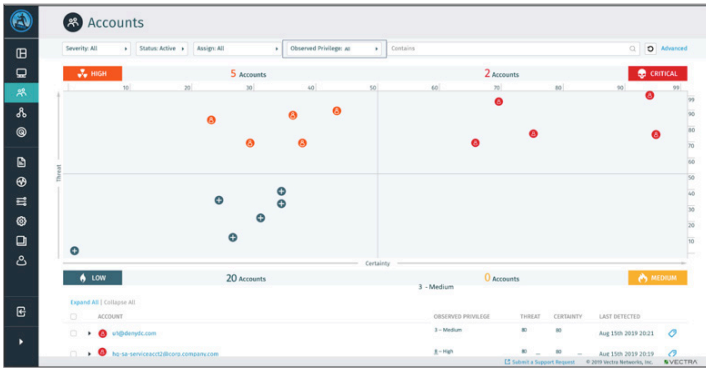
At the heart of the Vectra AI Platform are detection models that identify subtle indications of environment-knowledgeable attacks, all while only surfacing real events and eliminating noise.

The supervised AI models used by the Vectra AI Platform allow for very quick detections even from day one without any deployment delays. Vectra AI currently leverages more than 90% of the MITRE ATT&CK framework, and is constantly including more with every software release.

Privileged Access Analytics

A fundamental part of the engine that enables threat scoring is called Privileged Access Analytics (PAA). Rather than relying on the granted privilege of an entity or being agnostic to privilege, PAA focuses on how entities are actually utilizing their privileges within the network. This is known as observed privilege.

This viewpoint is similar to how attackers observe or infer the interactions between entities. In order to succeed, it is imperative that defenders think in a similar fashion as their adversaries.



Vectra AI Detect identifies and prioritizes all accounts that indicate anomalous behaviors

PAA starts by grouping observed identities into groups based on similarity. This grouping is the baseline for incurring what constitutes as normal and abnormal access patterns. Vectra AI then applies further models to detect common access attacks.

An account, like a domain admin for example, may have the rights to access any system within the entire network. However, it is probably not accustomed to doing that, and therefore its observed privilege might be lower than that of a service account that is used to deploying software updates onto thousands of systems on the network.

Conclusion


With the Vectra AI Platform, security analysts now have the tools needed to identity and prioritize observed privilege across all of their networks.

About Vectra AI

Vectra AI is the leader in AI security and observability. Vectra AI delivers organizations real-time visibility into their network, clear insight into which behaviors matter, and the ability to act before risk becomes impact. By connecting on-premises, multi-cloud, identity, SaaS, edge, and IoT/OT infrastructure, Vectra AI helps organizations reduce exposure, accelerate detection and response, and automate security operations with AI. With over a decade of AI and ML innovation, 39 patents and a Leader in the 2025 and 2026 Gartner Magic Quadrant for Network Detection and Response, Vectra AI empowers security teams to stay ahead of emerging AI powered attacks, increase operational efficiency, and prove resilient in an increasingly complex, AI-driven world. For more information, visit www.vectra.ai.

PAA is integrated across the entire Vectra AI Platform, as well as through APIs. In Vectra AI Detect, security professionals can find PAA under the Accounts Tab. This page surfaces all accounts in which Vectra AI has detected anomalies. The accounts are scored the same way the Vectra AI Platform scores hosts, on two axis for certainty and risk.

This is a very powerful way to give insight into what privileged identities exist on a network at any given time. Clicking on an account gives further context to an identity showcasing the observed privilege level fully enriched with Active Directory context, as well as the individual detections associated with that account for further investigation.



Vectra AI Stream features the same scalable security-enriched identity metadata in a Zeek-compatible format to feed custom detection and response tools