The Vectra AI Platform Overview

The Vectra AI Platform delivers AI-native security and observability across modern networks that span on-premises, hybrid, and multi-cloud environments. By applying Vectra AI's patented AI to security telemetry, the platform continuously observes, correlates, and prioritizes attacker behavior across domains so security teams can stop active threats. It also surfaces exposure across identities, assets, and environments so teams can reduce risk before impact.

The platform ingests and correlates telemetry across the modern network, including:

- Data Center (On-premises and cloud)
- Multi-cloud (AWS, Azure, GCP, OCI, IBM Cloud)
- Identity (Local network, Active Directory, Entra ID, Microsoft 365, Copilot for M365, AWS, Azure)
- SaaS (Microsoft 365, Copilot for M365)
- SASE (Zscaler, Netskope)
- IoT/ OT
- AI agents

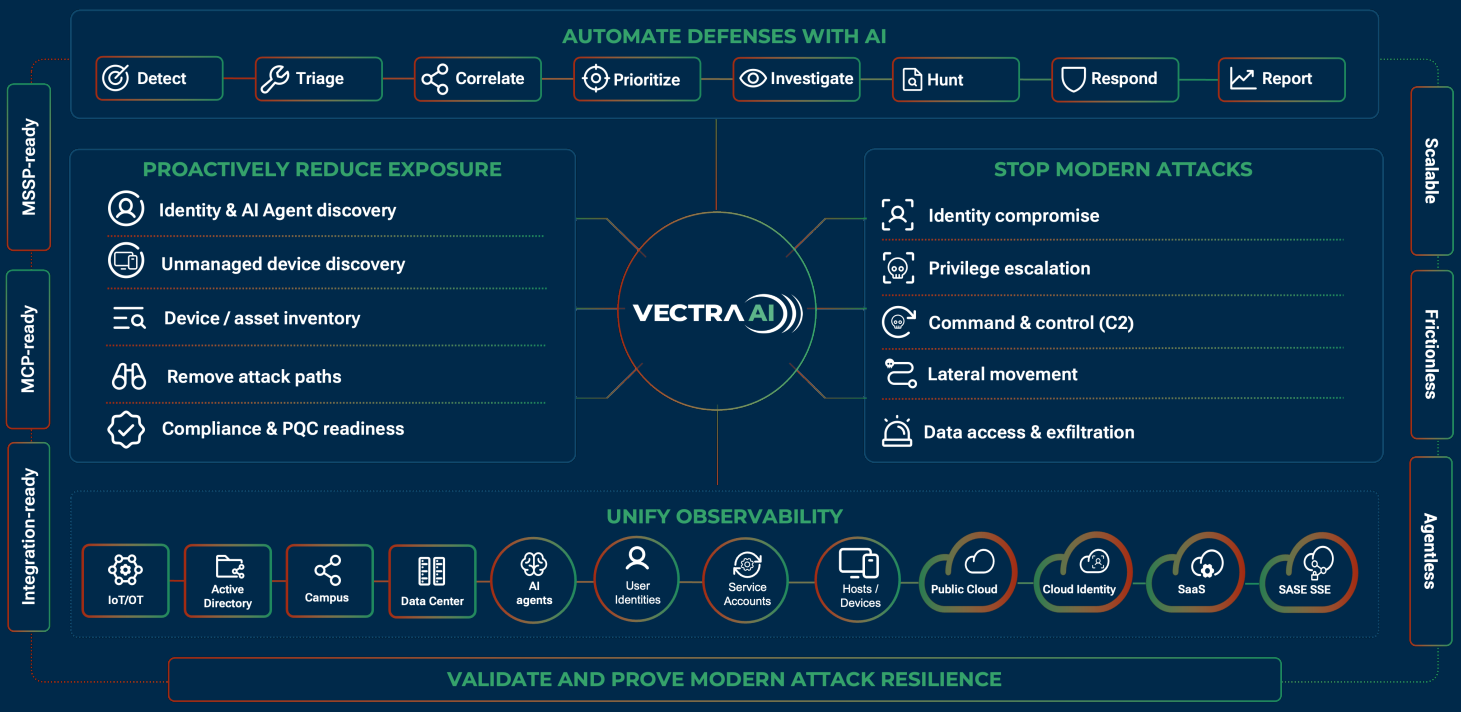

Table of Contents

Stop modern attacks

The Vectra AI Platform gives security teams a cross-domain view of attacker behavior across network, cloud, SaaS, and identity environments. Behavioral AI detects threats that EDR, SIEM, and native tools can miss, including account compromise, privilege abuse, lateral movement, command and control, and data access, so teams can investigate faster and stop attacks sooner.

DETECT

Behavioral-based detection

Vectra AI detects attacker behavior across the cyber kill chain by applying behavioral AI to network, identity, and cloud activity. By analyzing behavior instead of relying only on known indicators, Vectra AI can identify techniques such as reconnaissance, lateral movement, C2, and privilege abuse, including activity associated with zero-day exploits and AI-driven attacks.

Entity-centric detection model

Vectra AI groups detections by entity (e.g. device and identity) to consolidate related activity into one cross-domain investigation context.

Unified telemetry and cross-domain coverage and context

Vectra AI normalizes network, cloud, SaaS, and identity telemetry into a unified entity model. It links detections, authentication events, resource access, and communications to hosts and accounts so analysts can trace behavior across hybrid and multi-cloud environments.

Coverage	Data Sources	Detection Capabilities
Network	• Network metadata for east-west and north-south traffic across physical, virtual, and cloud environments • AI enhanced sub-second session data, flow records, and packet-derived metadata • Host-to-host communication patterns • Encrypted and unencrypted traffic without requiring decryption	• Detects lateral movement, command and control (C2), remote execution, reconnaissance activity, credential misuse, and privilege abuse • Uses deep machine learning and behavioral analytics to analyze large-scale network telemetry • Detects over 90% of MITRE ATT&CK techniques in network environments

Coverage	Data Sources	Detection Capabilities
Multi-Cloud	• Cloud control plane logs (e.g. AWS API activity) • Identity and access management (IAM) events, role assumption, and credential usage • Resource configuration and usage telemetry • Cloud flow logs, DNS logs, and network telemetry across AWS, Azure, GCP, Oracle Cloud, and IBM Cloud • Cloud provider metadata and 3rd party enrichment (e.g. security platforms, endpoint tools)	• Detects account compromise, privilege escalation, excessive permissions, role abuse, lateral movement, trust-boundary violations, hybrid attack sequences, reconnaissance, brute force, anomalous traffic, and data exfiltration • Identifies communication with suspicious external infrastructure using IP reputation and threat intelligence Leverages correlated multi-cloud telemetry (e.g. flow, DNS, control plane, and identity signals) to provide visibility into attacker activity across kill-chain and inter-cloud movement
SaaS	• Microsoft 365 audit logs • Microsoft Exchange, SharePoint, OneDrive, and Teams activity • Access and authentication events • Application access and API activity	• Detects account takeover, suspicious login and access patterns, malicious application behavior and OAuth abuse, and data access anomalies and exfiltration indicators
Identity	• Identity provider logs (Active Directory, Entra ID) • Authentication and authorization events • User and service account activity • Role assignments and permission changes • Token usage and session activity	• Detects compromised credentials and account takeover, anomalous authentication patterns, privilege escalation, abuse of administrative roles, lateral movement via identity systems, federated identity abuse, token-based access, and identity persistence mechanisms

IDS signatures

Vectra AI uses a single sensor and detection workflow, and combines Suricata-based signatures with behavioral AI analytics. Signatures identify known exploits and CVEs, while behavioral AI adds entity and activity context for unknown, evasive, and post-compromise attacker behavior. Security teams can validate threats, correlate signature and behavioral detections, and extend coverage beyond traditional IDS/IPS workflows.

SIEM and data lake export

Vectra AI exports enriched metadata and detections to SIEMs and data lakes. Normalized telemetry supports custom analytics, existing workflows, and external investigation beyond the Vectra AI Platform.

Threat intelligence

Vectra AI combines threat intelligence, signatures, and behavioral AI to detect known and unknown threats. Threat intel and signatures identify known malicious infrastructure, exploits, CVEs, and indicators, while behavioral AI detects evasive, zero-day, and AI-driven attacker behavior.

PRIORITIZE

AI-driven triage

- Evaluation of active detections for context, relationships between events, and commonalities across entities
- Automatically separates malicious from benign activity
- Reduces alert volume by filtering out false positives and weak indicators

AI-driven prioritization

- Correlates detections across domains and scores and ranks threats based on attacker profile, MITRE ATT&CK, severity, business impact, attack progression, attack velocity, data volume, breadth of attack, and rarity
- Supports custom prioritization based on critical assets and entity grouping

AI-driven correlation

- Correlates telemetry and detections from network, cloud, SaaS, identity, and authentication events
- Tracks devices across IP changes
- Normalizes identities across Active Directory, Entra ID, AWS, and SaaS
- Links users, hosts, accounts, and services
- Tracks attack progression across environments and stages
- Consolidates fragmented alerts into cohesive incidents

INVESTIGATE

Investigation contextualization and visualization

- Unifies cross-domain data into one investigation view with entity, privilege, activity, and risk context
- Attack graphs: visualize multi-domain activities, identify patient zero and reveal blast radius
- Attack timelines: show sequence and progression of attacker activity
- Attack flows: represent attacker movement across environments

Instant investigations

- Pre-built, context aware investigation information tied to detections and entities
- One-click access to relevant data scoped to time, entity, and activity
- Automatic correlation across network, identity, cloud, and SaaS activity without manual query building

AI-assisted investigation

- Converts natural-language questions into context-rich answers and recommended next steps
- Traces activity across network, identity, and cloud - from the first sign of compromise to lateral movement
- Generates consolidated threat summaries and attack timelines

Enhanced workflows through API and integrations

- Provides API access to detection, entity, and investigation context for downstream SOC workflows
- Sends correlated detections, related entities, severity, urgency, and supporting metadata to SIEM, SOAR, ITSM and Agentic SOC tools

AI-enhanced metadata

- Collects and normalizes metadata across network, identity, cloud, and SaaS environments, including connection activity, protocols (RDP, SMB, RPC), authentication events (Kerberos, NTLM), DNS, DHCP, and cloud control plane logs
- Enriches telemetry with host and identity context to track activity beyond IP addresses and across attack chains
- Provides high-fidelity, structured data (e.g. session characteristics, beaconing patterns, JA3/JA4 fingerprints) to support detection, investigation, and threat hunting without reliance on raw packet capture
- Enables efficient, large-scale analysis by transforming raw data into lightweight, searchable metadata, allowing teams to reconstruct activity, perform retrospective investigations, and validate exposure across environments
- Stores enriched network metadata for retrospective hunting, allowing teams to query historical host, identity, and communication activity, identify IOCs, and reconstruct attack timelines without external data sources

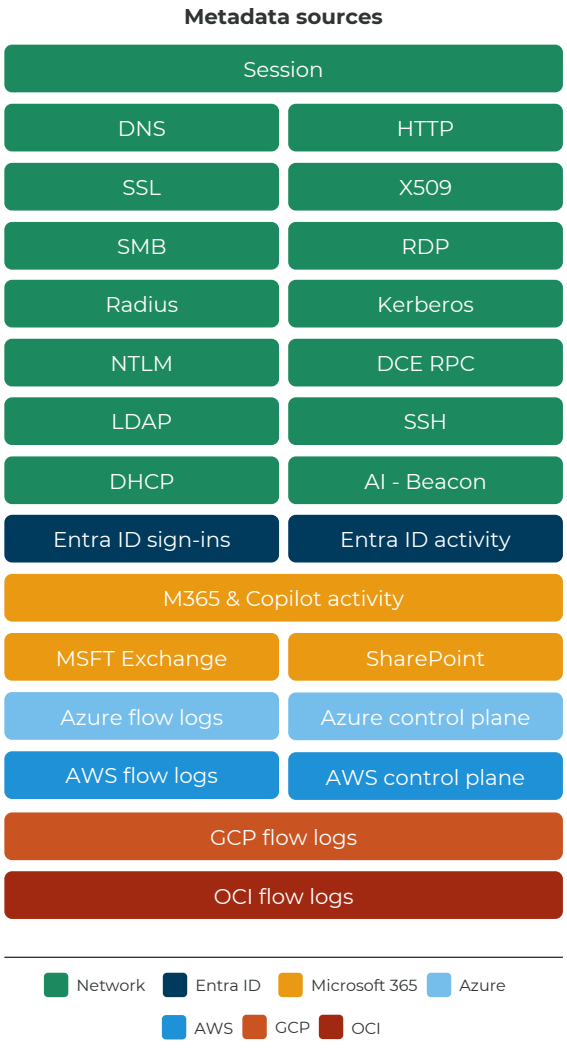

How Vectra AI enhances metadata

AI-stitching for network
Track machine names across IP changes

AI-stitching for cloud
Track cloud identities across log types, ObjectID, and role changes

AI-beacon data
Hunt and find risky call-backs

AI-privilege levels
Search for privilege hosts, services, and accounts based on AI graph analysis

JA4+
Track and hunt down data fingerprints

RESPOND

Native response actions (automatic or manual)

Response actions can be executed manually or driven by AI with customer-configured thresholds, including threat severity and entity importance.

- Device lockdown: isolate endpoints
- Identity lockdown: lock compromised accounts and trigger password reset
- Traffic lockdown: trigger containment workflows

Integration-based response

Vectra AI integrates with EDR, SOAR, and ITSM tools for manual or automated response and one-click pivots.

Managed response

Vectra AI managed services execute response actions through integrated EDR tools, including Microsoft Defender, CrowdStrike, and SentinelOne, and custom third-party integrations.

HUNT

AI-assisted hunt

- Translates natural language prompts into multi-step, structured hunts across cloud, identity, SaaS, and network telemetry
- Hunts for threat actor activity, validates exposure to new CVEs, identify weak configurations, and explore risk across domains
- Quickly understand baseline activity, identify anomalies, and proactively uncover risks or compliance gaps

5-minute threat hunts

- Pre-built hunting scenarios that are backed by research and target emerging threats, vulnerabilities, and risky behaviors
- Guided workflows with embedded context, including what to look for, why it matters, and recommended next steps
- One-click execution of prebuilt queries across unified modern network telemetry without manual query construction

Saved queries

- Store and reuse structured queries across multi-domain telemetry
- Reduce time to investigate by eliminating the need to rebuild complex queries

MANAGED SERVICES

Vectra MDR services provide 24x7x365 monitoring, analysis, exposure management, guided response, and collaboration to augment SOC operations.

Proactively reduce exposure

The Vectra AI Platform is a cyber risk decision engine that determines when to act, where to act, and where it matters. It provides continuous, data-driven insights that enable organizations to reduce exposure, demonstrate compliance, communicate risk, and optimize SOC performance.

Asset inventory

- Maintains a continuously updated inventory of: identities (non-human and human) and service accounts, network devices, and cloud resources
- Correlates assets across network, cloud, SaaS, and identity domains
- Associates assets with behavioral context and activity patterns
- Enables tracking of entity relationships, dependencies, and communication paths
- Provides the foundational data layer for detection, investigation, and observability insights

Observability insights

Vectra AI's observability insights use normalized telemetry and asset inventory for domain-specific visibility.

Threat surface dashboards
These dashboards provide visibility into exposure, activity, and risk across core domains.

Threat summary

Aggregates threats and activity across domains to highlight prioritized entities, attack activity, trends, risk posture, and active threats.

Azure observability

Provides visibility into Azure infrastructure activity and usage, surfaces suspicious behavior across compute, storage, and services, and correlates Azure activity with identity and network signals.

Entra ID observability

Provides visibility into unusual identity activity and authentication events, including stale accounts, excessive administrative privileges, accounts without MFA, and legacy authentication usage to support least privilege and identity governance reviews.

Network observability

Visualizes network activity, communication patterns, and traffic flows while surfacing risky protocols and monitoring privileged human and non-human identities.

Network map

Provides a graphical view of sensors, observed IPs, subnets, hosts, and connections with drill-down navigation and one-click pivot to connection-level SQL search.

PQC readiness

Identifies cryptographic protocols and algorithms in use, highlights potential exposure to quantum-vulnerable encryption, and helps prioritize cryptographic modernization.

AI readiness dashboards

These dashboards provide visibility into the use, risk, and governance of AI technologies within the enterprise.

AI observability

Detects unauthorized or unmanaged AI tool usage, identifies access to external AI services and APIs, and surfaces potential data exposure risk from unsanctioned AI activity.

Agentic AI estate

Identifies AI agents and automated workflows interacting with enterprise systems, tracks agent behavior, permissions, and access patterns, and surfaces risk related to automation, privilege use, and system interactions.

Copilot for M365

Provides visibility into Microsoft Copilot usage within M365, tracks how Copilot agents interact with enterprise data, and surfaces identity and access risks related to Copilot agents.

Exposure management findings

- Identifies exposed attack paths across identity, network, cloud, multi-cloud, and SaaS environments, including risky access paths, excessive permissions, misconfigurations, weak security hygiene, and unmanaged assets.
- Maps exposure findings to entities such as users, accounts, hosts, services, and cloud resources to show where risk exists and what could be impacted.
- Correlates exposed attack paths with observed attacker behavior to prioritize exploitable risk, validate control gaps, and determine where an attacker could move next.
- Continuously updates findings based on changes in assets, identities, permissions, connectivity, configuration, and observed activity.

Compliance and governance support

- Generates audit-ready reports based on detection trends, response actions, exposure reduction, and operational metrics
- Supports alignment with regulatory frameworks and recommends controls for detection and posture gaps

Gap exposure assessments

Vectra AI offers red team, blue team, and purple team workshops and protected open-source emulation attack tools for testing identity, cloud, and network environments.

Other platform capabilities

Executive-level reporting

- Aggregates prioritized threats, identity risks, and detection trends into leadership-ready insights
- Highlights concentration of risk across environments, entities, and attack activity
- Translates technical signals into business-relevant metrics for clear communication of security posture and improvements

SOC optimization with data-driven insights

- Analyzes detection trends, MITRE ATT&CK-aligned behaviors, and recurring entities to identify patterns and inefficiencies
- Surfaces insights into triage accuracy, investigation efficiency, and MTTx (e.g. MTTD, MTTR)
- Enables continuous refinement of SOC workflows, improving prioritization, response effectiveness, and operational performance

Services

Vectra AI services include technical support, implementation, configuration, and training.

Technical support
Vectra AI's technical support offers customers a combination of outstanding customer support and best-in-class technical ability and engineering agility.

Professional services
Vectra AI's professional services include implementation services and customer training. These services come as available add-ons to Vectra AI products and are meant to assist in the success of our customers on the Vectra AI Platform with hands-on expertise from our professional services and support teams.

Integrations

Vectra AI Platform integrates with a broad ecosystem of security and IT tools, including SIEM, SOAR, EDR, ITSM, and network infrastructure solutions. This interoperability allows organizations to incorporate Vectra AI into existing architectures while maintaining consistent workflows across detection, investigation, and response processes.

Packet broker	cPacket, ixia, endace Record. Respond., KEYSIGHT, Gigamon®
Networking	cisco, CHECK POINT, Microsoft Hyper-V, Microsoft Active Directory, JUNIPER NETWORKS, paloalto NETWORKS, zscaler, NUTANIX, NOZOMI NETWORKS, FORTINET, superna, KVM

Cloud	aws, Google Cloud, Microsoft Azure, Microsoft 365, Microsoft Entra ID, IBM Cloud, ORACLE CLOUD
Endpoint	CROWDSTRIKE, SentinelOne, FIREEYE, Microsoft Defender, VMware Carbon Black™, cybereason
SIEM/SOAR/ITSM	splunk> a cisco company, Chronicle, CROWDSTRIKE Falcon LogScale, Microsoft Sentinel, Radar, elastic, CORTEX XSOAR BY PALO ALTO NETWORKS, SWIMLANE, servicenow
SASE/SSE	zscaler, netskope

Deployment and architecture

The Vectra AI Platform combines flexible deployment across physical, virtual, and cloud environments with a scalable sensor and processing architecture that delivers continuous visibility across network, identity, and cloud activity.

Deployment

Deploy agentless as on-premises, SaaS, or hybrid and get actionable attack signal in days to hours across the network, and hours to minutes for identity and cloud.

Simplified deployment steps:

- Vectra AI instance is created.
- Welcome email is sent, and admin configures roles, users and any non-network data sources.
- Vectra Brain appliance deployed and connected to Vectra AI instance.
- Vectra Sensors are paired to the Vectra Brain and network traffic is directed to Sensors.

Architecture

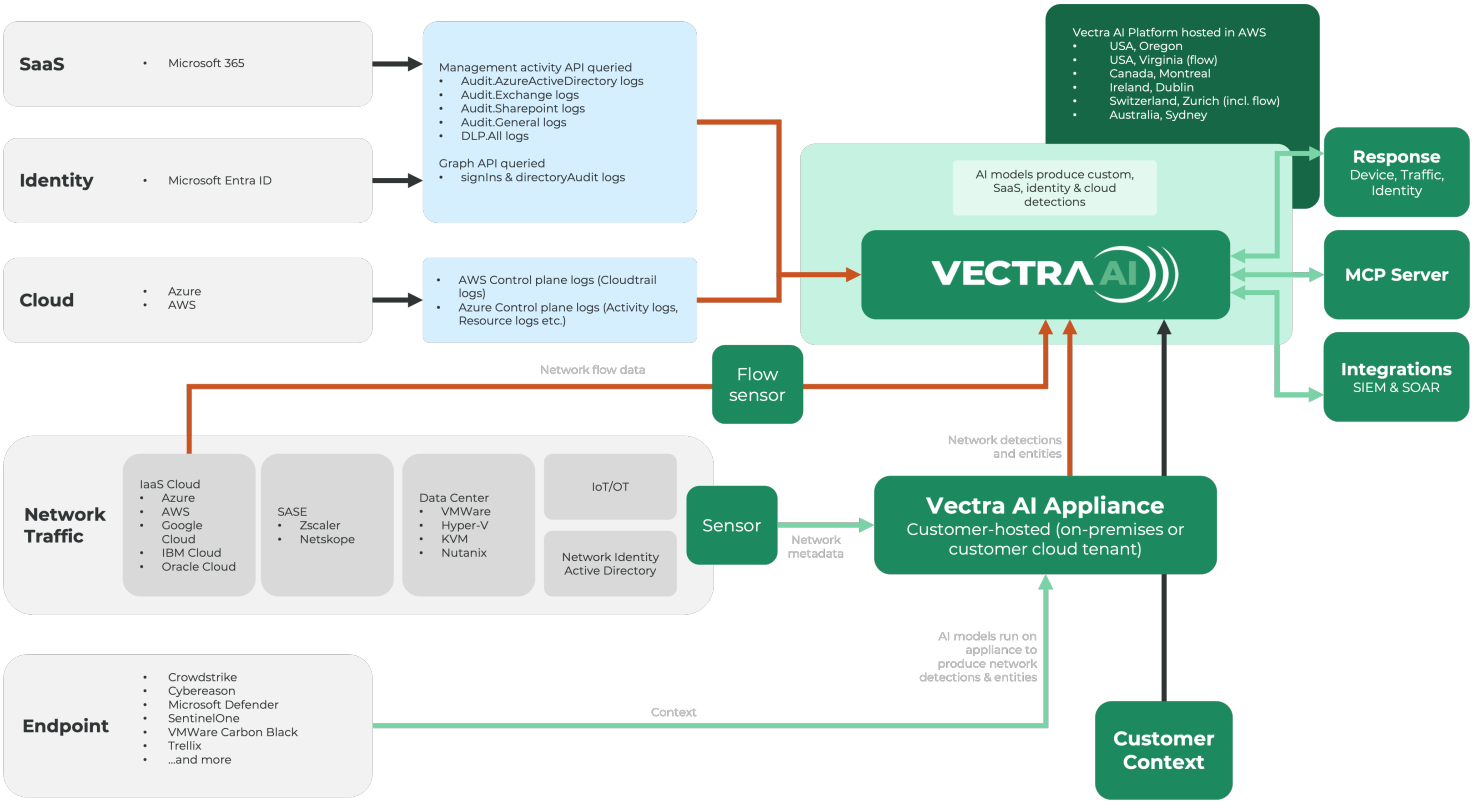

Appliance and sensor specifications

- Supports deployment across physical appliances, virtual sensors, and cloud environments (AWS, Azure, GCP), with traffic ingestion via SPAN ports, TAPs, packet brokers, and native cloud mirroring (e.g., VPC Traffic Mirroring, VTAP)
- Uses a scalable sensor and processing architecture to capture network activity and convert it into enriched metadata for analysis
- Separates data collection (Sensors) from processing (Brain), where sensors passively capture traffic and forward metadata to centralized or distributed processing nodes that run detection models and integrate with the Vectra AI cloud for investigation and response workflows

For more information about Vectra AI's appliance and sensor specifications, see here.

Summary

The Vectra AI Platform delivers AI-native security and observability to stop active threats and proactively protect the modern network. By correlating telemetry across network, cloud, SaaS, and identity domains, it delivers high-confidence detection of attacker behavior, reduces noise through automated triage and prioritization, and enables fast, context-rich investigation and response. The platform also provides continuous observability into assets, exposure, operational performance, and risk, supporting compliance, executive reporting, and ongoing SOC optimization. Delivered as a scalable SaaS solution with sensor-based deployment options and broad ecosystem integrations, Vectra AI helps security teams detect, investigate, hunt, and respond across hybrid and multi-cloud environments.

See the Vectra AI Platform in action

Want to chat with a Vectra AI Specialist?

Contact our team