VECTRA AI

REPORT

2026 State of Threat Exposure Management Report

Executive summary

Exposure is becoming increasingly dynamic.

Modern enterprise environments are changing faster than most security programs can accurately track. Assets continuously appear and disappear. Identities increasingly operate without direct human involvement. AI-driven systems, automation pipelines, unmanaged devices, cloud workloads, and legacy infrastructure now coexist across highly interconnected environments.

At the same time, security times remain heavily dependent on fragmented visibility sources designed for environments that were more static and easier to inventory with traditional security tools.

To better understand how this gap affects modern exposure management, Vectra AI analyzed customer environments across asset activity, AI-agent presence, and observed exposure conditions.

What we found was consistent across environments:

1. Asset change is continuous.
2. The number of unmanaged devices is significant.
3. AI agents are already materially present.
4. Exposure conditions remain widespread.
5. Visibility gaps continue to create operational uncertainty.

The findings suggest that exposure management is no longer primarily a vulnerability-management problem. It is increasingly a problem of continuously understanding what is operating, communicating, and creating exposed attack paths across dynamic environments.

Section 1: Observations

OBSERVATION 1:

Asset change was universal across observed environments

Over a 14-day period, every analyzed customer environment experienced newly observed hosts.

The findings suggest that modern enterprise environments are continuously changing, creating persistent inventory drift and operational uncertainty.

14-DAY ASSET-CHANGE OBSERVATIONS

100%
of environments saw new **devices appear**

90%
of environments saw new **device roles emerge**

83%
of environments saw new **device types introduced**

The findings do not suggest that every newly observed host is inherently risky. Instead, they highlight how frequently environments introduce uncertainty around how the following could create exposure or become part of a potential attack path:

- Ownership
- Role
- Management status
- EDR coverage
- Segmentation
- Communication behavior

What this means:

Modern environments are continuously changing faster than traditional inventories can accurately track. Security and infrastructure teams can no longer assume they have a complete understanding of what is operating, managed, segmented, or protected at any given moment.

OBSERVATION 2:

EDR visibility does not extend to the entire environment

Across observed customer environments, on average, **more than 30% of devices are unmanaged** because endpoint agents could not be deployed on them.

These devices include a mix of:

- IoT and OT systems
- Network infrastructure
- Printers and embedded devices
- Legacy systems
- Third-party and contractor-owned assets
- Specialized workloads

This finding suggests that agent-based security visibility alone cannot provide a complete view of modern enterprise environments. As organizations adopt more specialized, distributed, and non-traditional assets, unmanaged devices continue to represent a meaningful portion of the attack surface and create paths that endpoint tools cannot see.

What this means:

Organizations cannot assume their endpoint tools see everything operating in the environment. A significant percentage of assets often remain outside the reach of agent-based controls, creating blind spots that obscure exposure, weaken visibility, and leave potential attack paths unmonitored.

OBSERVATION 3:

AI-agents and non-human actors are already material

Among customer environments with monitored AI-agent activity over a 90-day period, AI agents and non-human actors were present at a scale large enough to materially affect security operations and exposure management.

90-DAY AI-AGENT OBSERVATION

Typical environment has roughly one AI agent per device:

1.17 **AI agents per device**

Extreme case: up to

96 **AI agents on a single device**

35% of environments had **more AI agents than devices**

The findings suggest that security teams increasingly need visibility beyond traditional assets and user identities. Modern environments now include:

- AI agents
- Service accounts
- Automation pipelines
- APIs
- Non-human identities
- SaaS-driven workflows

These entities authenticate, move data, trigger downstream actions, and interact across systems continuously, creating new access and communication paths that security teams need to observe and validate.

What this means:

Security teams now need visibility beyond traditional users and devices. AI agents, automation workflows, service accounts, and non-human identities are becoming active operational actors that new access paths, communication workflows, and potential attack paths across the environment.

OBSERVATION 4:

Attacker-relevant exposure conditions remained widespread

Across analyzed customer environments over a 30-day period, risky exposure conditions were consistently present.

30-DAY EXPOSURE FINDINGS OBSERVATIONS

98% **of environments**

had at least one attacker-relevant exposure condition present

63% **of environments**

showed exposure across multiple risk themes, including legacy protocols, weak cryptography, credential exposure, and exposed remote access

The most commonly observed conditions were not obscure edge cases. They reflected familiar operational weaknesses that attackers use to gain access, move laterally, steal credentials, weaken encryption, or expand control once inside an environment.

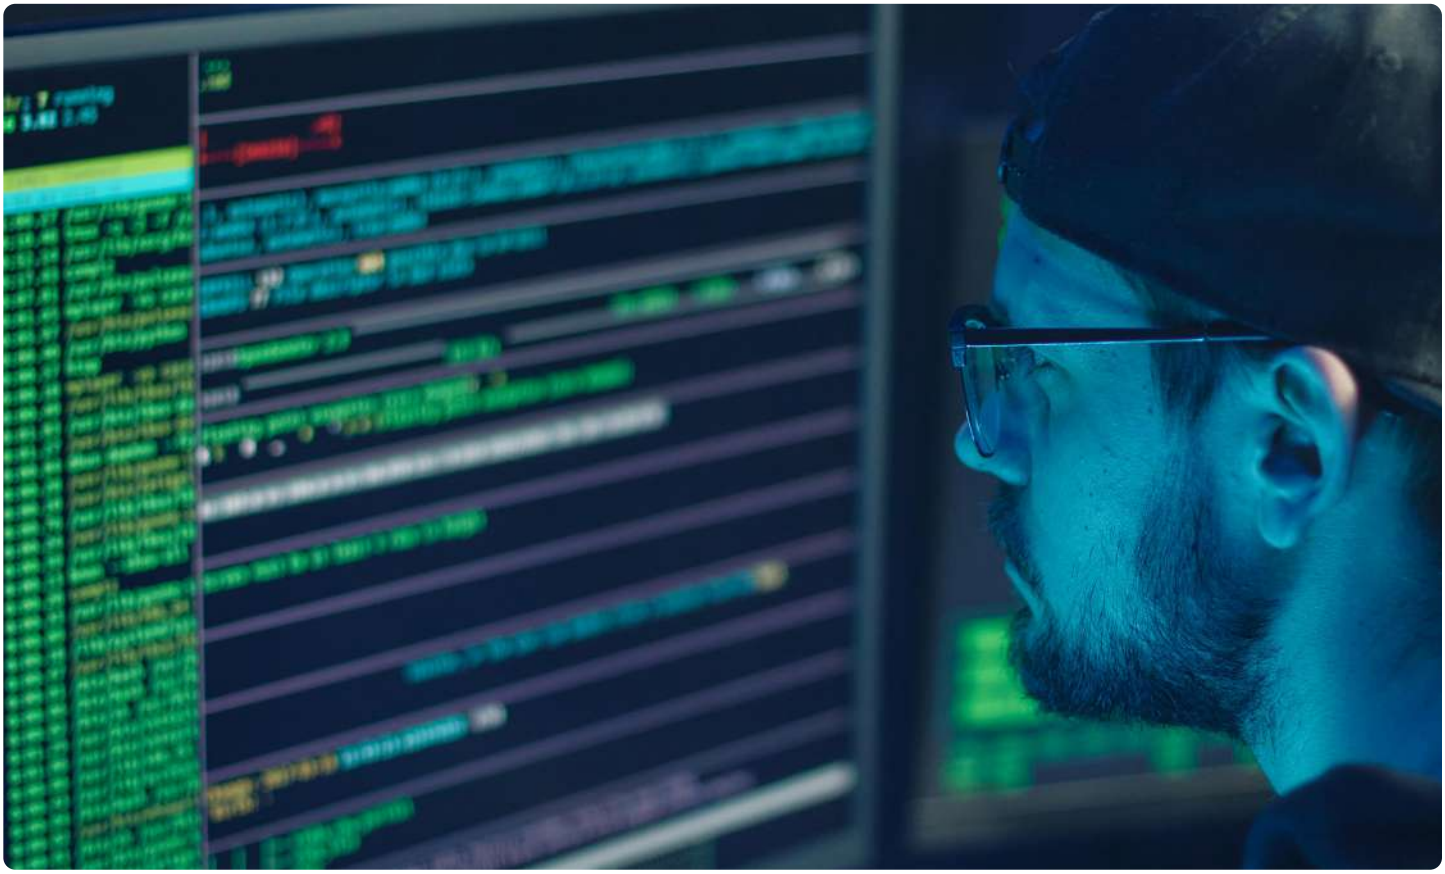

Most commonly observed attacker-relevant exposure conditions

OBSERVED CONDITION	SHARE OF ANALYZED CUSTOMER ENVIRONMENTS	WHY IT MATTERS
Deprecated TLS client	**96%**	Weak or outdated cryptography can increase attack surface and signal unmanaged infrastructure.
Expired certificates	**91%**	Certificate issues can disrupt operations, weaken trust, and expose gaps in cryptographic hygiene.
NetBIOS	**86%**	Legacy name-resolution protocols can expose internal information attackers use for discovery and lateral movement.
Plaintext passwords	**85%**	Exposed credentials can accelerate compromise, privilege escalation, and lateral movement.
Deprecated TLS server	**82%**	Older TLS versions increase misconfiguration risk and may expose systems to weaker encryption.
Certificates expiring soon	**82%**	Upcoming certificate expirations create operational risk and reveal systems that may not be actively managed.
Credential files in SMB	**81%**	Credential files in shared locations can give attackers reusable access to additional systems.
FTP	**73%**	Unencrypted file transfer can expose sensitive data and credentials in transit.
SMBv1 server	**61%**	SMBv1 is a legacy protocol commonly associated with lateral movement and ransomware risk.
Telnet	**57%**	Telnet transmits data without encryption, making credentials and administrative access easier to intercept.
SMBv1 client	**55%**	SMBv1 client activity can indicate continued dependence on insecure legacy communication.
Exposed RDP	**46%**	Exposed remote access can provide attackers with a direct entry point or persistence path.

The findings suggest that attacker-relevant exposure is widespread, and security teams need context to determine which conditions create the greatest risk based on asset sensitivity, segmentation, identity access, communication behavior, and operational role.

What this means:

Legacy protocols, weak cryptographic posture, exposure services, and credential-related risks continue to persist across modern environments despite ongoing modernization efforts. Security teams need continuous operational context to understand which conditions are merely present, which ones are exploitable, and which ones should be remediated first.

Section 2: Key Takeaways

1

Why traditional visibility models are struggling

The observations across all three areas point toward a broader issue: modern environments evolve faster than traditional visibility models can accurately represent. Most security programs still rely heavily on periodic scans, static inventories, siloed telemetry, and point-in-time assessment.

However, the observed environments behaved continuously and dynamically. Assets changed constantly. Actors expanded beyond human users. Exposure conditions persisted across interconnected systems, creating paths that static visibility models struggle to represent.

This creates a growing gap between what organizations believe exist and what is actively operating.

2

Exposure is becoming behavioral, not static

The findings suggest that exposure management is increasingly less about isolated vulnerabilities and more about understanding behavior, relationships, and operational context.

Exposure now emerges through identity usage, communication paths, delegated trust, unmanaged infrastructure, AI-driven automation, and cryptographic dependencies – the same relationships that can become attack paths when left unvalidated. This changes how organizations must think about risk prioritization. For example, a vulnerability without access or communication may present limited operational risk. However, a trusted identity moving across sensitive systems may present significantly greater risk exposure because it can create a path to sensitive systems even without a traditional CVE.

3

The shift toward continuous, evidence-based exposure management

The observations suggest organizations increasingly need exposure management approaches grounded in continuous operational evidence rather than periodic interpretation.

That includes:

- Continuously updated asset understanding
- Visibility into non-human actors
- Communication-aware exposure context
- Validation of boundaries and segmentation
- Prioritization based on observed behavior

The challenge is no longer simply collecting more telemetry. It is developing a continuously updated understanding of what is operating, what is changing, what is interacting, and what creates exploitable attack paths.

Conclusion

Modern enterprise environments are dynamic by default.

Assets continuously change. AI agents and automation expand operational complexity. Legacy exposure conditions persist. Trust relationships evolve faster than static systems can accurately model, creating exposed paths that security teams need to continuously validate.

The organizations best positioned to reduce exposure going forward will likely be those capable of continuously observing how their environments behave, not simply how they are designed or documented, so they can identify and reduce exposed attack paths before attackers exploit them.

Learn more about Vectra AI's **Network Threat Exposure Management & Posture Improvement**

Appendix

Methodology

This report is based on aggregated observations across customer environments monitored by Vectra AI.

Observations were collected across anonymized customer environments over 14-day, 30-day, and 90-day observation windows during the study period.

The data reflects network-observed activity across hybrid enterprise environments, including combinations of on-premises infrastructure, cloud workloads, SaaS environments, identity systems, and unmanaged and managed devices.

Asset observations were derived from network-observed communication behavior. "New device types" represent newly observed MAC vendors and serve as a directional proxy for device diversity and inventory drift rather than definitive device fingerprinting.

AI agent observations reflect monitored non-human and AI-driven entities identified through observed activity patterns, service interactions, and automation behavior within participating environments.

Exposure findings reflect observed operational and security conditions identified through network-observed behavior, protocol usage, certificate activity, and communication analysis.

All customer data was anonymized and aggregated prior to analysis.

About Vectra AI

Vectra AI is the leader in AI-native 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 please contact us:
Email: info@vectra.ai | vectra.ai