+

Vectra AI Platform – 9.2 Release Notes

The Vectra® X-series appliances, B-series appliances, S-series sensors, VM and Cloud Brain, VM and Cloud Sensors, are scheduled to be updated to Vectra Network software release version 9.2. The 9.2 release includes expanding GCP brain offerings, the Executive Overview Report, adding groups support for the QUX API, and AI-Triage for Azure Cloud and AWS Cloud Detections.

9.2 will have the following release schedule:

- **Customers with Remote Support Enabled:** Customers who have remote support enabled will receive the update on or after June 26th, 2025.
  - You can check if you have remote support enabled under Settings > General with Remote Support set to Enabled.
  - If you plan to enable or disable Remote Support in the near future, please reach out to Support to confirm if you will receive or skip the upgrade.
- **Customers Connected to Updater**: Customers who do not have remote support enabled but are connected to Updater will receive updates on or after July 7th, 2025.
  - You can check if you are connected to Updater under Data Source > Brain-Setup > Proxy & Status and see that Updater Destination shows as connected, while Remote Support shows disabled.
- **All Other Customers***: Will be able to download the update on or after July 7th, 2025.
  - *Note: This does not impact customers that have requested they be pinned to a specific release from support.

Platform

Expanding GCP Brain Offerings — Network

Starting in 9.2, Vectra is introducing additional Brain offerings hosted in Google Cloud Platform, or GCP. The new GCP Brains are capable of handling 5Gb/s and 15Gb/s and support all the same features as other Cloud/Virtual/Hardware Brains. Please see the GCP Brain Deployment Guide for details.

Executive Overview Report — Network

Starting in 9.2, Vectra is introducing the Executive Overview report on the Vectra AI Platform. This report is catered to CISOs and security executives who need to bring high-level metrics to their board or executive-level meetings. Metrics include noise to signal tunnel, attack trends, and more. This report allows executives to make strategic decisions and evaluate how Vectra reduces security breach risk for their organization.

Added Group Support Added for v2.x. QUX API — Network

Starting in 9.2, Vectra supports getting group members from the /groups endpoint. For more information see: https://support.vectra.ai/vectra/article/KB-VS-1638

AI-Triage for AWS Cloud and Azure Cloud Detections — Cloud Coverage

Vectra AI has introduced AI Triage, its proprietary agentic AI solution to its AWS and Azure coverage portfolios. AI-Triage now auto-investigates AWS Cloud and Azure Cloud alerts based on factors such as prevalence and threat profiles to filter benign activities in customers' environments. The impact of AI-Triage is a reduction in prioritized entities and corresponding investigation workloads for SOC analysts.

Detections

Suspect Protocol Activity: Internal Detections — Network

Vectra is expanding the coverage of the Suspect Protocol Activity detections. Now, Suspect Protocol Activity includes detections covering Internal Lateral/Recon attacks and supports LDAP, Kerberos, NTLM, and SMB protocols. This feature is off by default but can be customer enabled and is included as part of the standard Detect product line. For more information on SPA, please see https://support.vectra.ai/s/article/KB-VS-1793

Suspect Protocol Activity: Brute Force — Network

Vectra is expanding the coverage of the Suspect Protocol Activity detections. Now, SPA can detect brute force attempts over all protocols. This rule detects brute force attacks where an attacker attempts multiple authentication requests in a short period. Brute force attacks can target various protocols such as SMB, LDAP, FTP, RDP, SSH, and HTTP, and are often used by adversaries to gain unauthorized access to accounts.

New Detection: NTLM Relay Activity — Network

Vectra AI has introduced a new detection for NTLM Relay Activity. This enhances Vectra's visibility into lateral movement techniques used by attackers. This detection identifies attempts to exploit NTLM authentication by observing when an attacker queries one host and relays the captured authentication to another host—often as part of privilege escalation or domain compromise efforts.

New Detection: M365 Copilot Sensitive Data Discovery — M365/AAD

Vectra AI has introduced a new detection for discovery behaviors surrounding M365 CoPilot. The new M365 CoPilot Sensitive Data Discovery detection where a CoPilot session was leveraged by an identity to access file(s) that may contain sensitive information. This detection aims to surface threat actors that use an account in the environment to discover sensitive information.

New Detection Suite: AWS Bedrock Detections — AWS

Vectra AI has introduced four new detections to surface suspicious behaviors surrounding the use of AWS Bedrock, a fully managed service offered by AWS that simplifies building and deploying generative AI applications.

- **AWS Bedrock Logging Configuration Disabled**: This detection highlights instances where a principal was observed disabling prompt logging for AWS Bedrock at the regional level. Disabling prompt logging stops the capture of all prompt and response activity across AWS Bedrock models and may indicate an attempt to impair defenses or hide malicious usage.

- **AWS Bedrock Novel Model Enabled:** This detection identifies suspicious activity related to the enablement of an AWS Bedrock Model by an identity that has no prior history of performing such actions. It flags potential unauthorized access to generative AI services that may be security-sensitive and associated with high-cost.
- **AWS Suspicious Bedrock Activity:** This detection identifies suspicious activity related to the enablement and invocation of an AWS Bedrock Model by an identity that have no prior history of performing such actions. The combination of enablement followed by invocation of a model suggests an attacker is both testing and using the model, generating responses at the victim's expense.
- **AWS Bedrock Novel Enabled:** It detects every instance when an AWS Bedrock foundational model is enabled, as this action is uncommon and may have cost or security implications. This is an informational detection and does not contribute to scoring or prioritization of the entity. It is meant to be a security relevant insight and may not be deemed immediately suspicious.

Signal Enhancements

Significantly reduced benign prioritization alerts through improvements to Vectra's AI prioritization algorithm and detection updates. In some cases, customers may see up to 50% fewer prioritized host and account alerts—without sacrificing coverage for real threats.

- **Azure AD & M365:** Prioritization alerts for accounts with specific detections have been refined, reducing benign alerts while maintaining detection of modern attacks. Affected detections include M365 Suspicious Download Activity, which now incorporates Autonomous System Number (ASN) context and Azure AD Suspicious Scripting Engine, with improved parsing for user agents.

Rapid Release Improvements

The following improvements have been made to algorithms since the last software release cycle. Customers that are connected to Vectra's update service with Remote Support enabled have received these improvements. All other customers will be receiving the following improvements as part of 9.2:

- NDR-166: This release enhances DNS Tunnel detection by expanding coverage across all DNS response types, providing broader and more accurate threat detection.
- NDR-144: Improves C2 detections against techniques used by the Covenant C2 Framework.
- NDR-202: This release enhances the performance of the algorithm powering our Exfiltration detections, enabling faster threat identification.
- NDR-195: Improves HTTP detections against penetration techniques used by the Kali Linux Package Repository.
- NDR-221: Improves HTTP detections against suspicious usage of Windows Remote Management (WinRM), strengthening visibility into potential abuse of this protocol.
- NDR-232: Enhances Suspect HTTP Activity detections to account for proxy usage, improving detection accuracy in proxied environments.

Bug Fixes

TITAN-2301: Network Traffic Validation Report Download

Resolved an issue where the Traffic Validation Report failed to download when the report exceeded around 6MB.

CS-10241: Some Detections Missing in the GUI

Resolved an issue where customers would be alerted to detections via Syslog, but the detections would not appear in the Detection GUI. This issue was solved by ensuring the detection ID stays consistent throughout the entire pipeline.

CS-10220: SPA Detections Not Working with Southside Proxies

Resolved an issue where Suspect Protocol Activity Detections do not trigger when using proxies but does without. The issue was resolved by ensuring the list of proxies is on both brain and sensor.

CS-10145: Kerberoasting: Targeted Weak Cipher Responses PCAPs Not Downloadable

Resolved an issue where the PCAPs for the detection Kerberoasting: Targeted Weak Cipher Responses could not be downloaded from the UI. This issue has been resolved.

Appendix:

Will this upgrade perform a reboot of the Brain or Sensors?

No, a reboot is not required as part of the 9.2 update.