

Vehere Vision AI: Agentic AI for Alert Analysis in NDR

Executive Summary

Vision AI is an agentic AI capability built into Vehere's Network Detection and Response (NDR) platform. It reviews the alerts produced by the platform's machine-learning (ML/DNN) detection models and decides, for each one, whether it is likely a genuine threat (True Positive) or benign (False Positive). It then explains its reasoning, suppresses noise, proposes whitelist entries, enriches serious alerts with external threat intelligence and recommends response actions.

Vehere positions this in a way to cut alert fatigue and speed up triage for security operations center (SOC) teams, while keeping the analyst in the loop through visible reasoning and configurable confidence thresholds.

Key points at a glance

Scope: Designed for alerts from ML/DNN- based analysis of raw session traffic. It does not cover other alert types.

Core functions: Alert classification, noise reduction, automated whitelisting, threat enrichment and response support.

Integration: Outputs are designed to feed playbooks and SOAR workflows.

Control model: Analysts set confidence thresholds that decide which alerts are hidden or highlighted.

Dependency: The enrichment feature (Deep Threat Insights) requires internet connectivity from the deployment.

Overview

Vision AI continuously analyzes ML/DNN-generated alerts using behavioral analysis and historical context. For each alert it draws on:



Alert metadata



Protocol behavior



Session information



Contextual indicators

The result is shown directly in the platform's Alerts module and in the Alert Details panel, so analysts do not need a separate tool

Prerequisites and Availability

Requirement	Detail
Enablement	Vision AI is switched on through the platform's advanced configuration settings
Alert coverage	Available only for ML/DNN-based raw session traffic alerts
Connectivity	Deep Threat Insights needs internet connectivity in the deployment environment
Network access	A dedicated port must be open for Deep Threat Insights communication.

Buyers with air-gapped or tightly segmented networks should note that classification and noise reduction do not depend on internet access according to the vendor material, while enrichment does.

Core Capabilities



Alert Classification and Prioritization

Vision AI labels each alert as a True Positive or False Positive using traffic patterns, behavioral analysis and historical consistency. High- confidence True Positives are highlighted so analysts can triage them first, and results appear in the Alerts module alongside the original alert



Noise Reduction and Optimization

The system automatically suppresses low-confidence False Positives and filters repetitive, non actionable alerts. It applies tuning mechanisms to improve alert quality over time. The stated aim is to reduce alert fatigue and keep analysts focused on critical threats.



Automated Whitelisting

When the AI identifies recurring benign traffic, it can create a whitelist entry automatically. Entries may include source IP, destination IP, ports and protocols. This prevents analysts from repeatedly investigating the same false alerts and supports long-term tuning.



Automated Response Support

The agent gives investigation guidance and actionable recommendations and suggests mitigation and remediation steps. Its output is designed for playbook- driven operations and SOAR-ready workflows, with the stated goal of making response more consistent across the SOC.



Contextual Threat Enrichment

Vision AI uses external threat intelligence to give a high-level threat summary and key findings. It identifies attack vectors, suspicious protocol behavior and risks, so analysts can understand the wider scenario without researching it by hand

Severity Levels and Confidence Thresholds

Vision AI supports three severity levels (Low, High and Critical) and three confidence levels (Low, Medium and High). Two threshold settings let each organization decide how much AI output to act on

False Positive threshold: which False Positive alerts remain visible

Settings	Alerts shown
Low	Low-confidence False Positives only
Medium	Low-and medium-confidence False Positives
High	All False Positives
None	No False Positives shown

True Positive threshold: which True Positive alerts are highlighted in red

Settings	Alerts highlighted
Low	All True Positives
Medium	Medium-and high-confidence True Positives
High	High-confidence True Positives only
None	No highlighting

These controls are a notable strength. A cautious team can keep every False Positive visible while it builds trust in the model, then tighten the settings as confidence grows.

Analyst Experience

Enhanced Alert Details

The Alert Details panel shows the Vision AI classification and an expandable analysis. The analysis includes:


- Detailed reasoning for the classification
- Key indicators such as traffic patterns and protocol behavior
- Recommended actions for the analyst
- MITRE D3FEND defensive techniques where applicable

Deep Threat Insights

Deep Threat Insights adds a richer contextual analysis inside the same panel. It includes threat intelligence summaries and key findings and highlights risks, attack vectors and suspicious activity. The analyst triggers it with a Generate button, which becomes available only when Internet Search is enabled.

The design choice to keep both views inside the existing alert workflow reduces tool switching and keeps the reasoning attached to the alert for later review.

Business Value



Analyst Assessment

Strengths

- **Transparent output.** Each verdict carries reasoning, key indicators, recommended actions and a confidence rating
- **Adjustable trust.** Separate thresholds for False Positive visibility and True Positive highlighting let teams adopt the feature gradually.
- **Focused scope.** Tying the capability to ML/DNN session alerts keeps the problem well defined.
- **Closed-loop tuning.** Suppression and automated whitelisting reduce repeat noise without manual rule writing
- **Standards alignment.** MITRE D3FEND mapping and SOAR-ready output fit existing SOC practice.

Summary

Vision AI applies agentic AI to a well-understood NDR bottleneck: too many alerts, too little context. Its main strengths are explainable verdicts, adjustable confidence thresholds and in-workflow enrichment. The case for adoption will rest on measured evidence of accuracy and efficiency gains, plus clear safeguards around automated suppression and external data use. Analysts evaluating the product should request that evidence.