

Vehere Vision: Agentic AI for Network Detection & Response

Executive Summary

Security operations centers (SOCs) that rely on machine-learning-based Network Detection and Response (NDR) face two persistent problems: a high volume of alerts that must be triaged by hand, and a lack of immediate threat context when an alert looks serious. Vehere has addressed both by adding an agentic AI layer to its Vision NDR platform.

The layer consists of a single AI agent that delivers two capabilities:

Network Alert Analyzer (noise cancellation):

Classifies each ML-generated network alert as a True Positive or False Positive, explains the reasoning, and supports automated response and whitelist tuning

Deep Threat Insight: Turns a raw alert into a structured, context-rich threat intelligence brief within seconds, replacing manual research across external sources

The design keeps the analyst in control. The agent produces a verdict, a stated confidence level, the rationale behind it and recommended actions, all inside the alert-detail view the analyst already uses

Market Context and Problem Statement



Alert fatigue

ML detection models generate alerts at a volume that SOC teams struggle to triage manually. Manual triage delays response to genuine threats, and a high proportion of false positives wears down analyst attention.



Missing threat context


For more advanced threats, analysts need depth that the platform's own data did not provide. Without an integrated threat repository, analysts searched external sources by hand before they could act, which slowed triage.

These are well-recognized pain points across the NDR category, and the two capabilities below map directly onto them.


Solution Architecture

The Vision agentic architecture sits on top of the existing detection pipeline rather than replacing it.

Detection pipeline (existing)



Agentic layer (new): A single AI agent, covering noise cancellation and deep insight, reads from and writes to that database. It can also perform internet lookups to enrich its analysis. It produces three outputs:




All outputs are presented to analysts through the Vision web UI. Because the agent's reasoning is stored alongside the original alert, the context persists for later review and tuning.

Capability 1: Network Alert Analyzer (Noise Cancellation)

What it does. The agent evaluates each ML-generated alert using the alert description, session data and historical communication context. It then classifies the alert as a True Positive or False Positive.

Automated response. For real threats, the platform supports automated response actions such as next generation firewall (NGFW) control and TCP reset. For noise, the agent reduces repeat alerts through whitelisting and tuning.

What The Analyst Sees. The Alert-Detail View Shows:



Illustrative examples from the vendor demonstration

Scenario	Verdict	Basis for the verdict	MITRE tactic
Web application attack	True Positive, High confidence	Repeated HTTP requests carrying remote file inclusion parameters, the Nikto scanner user agent, an internal destination and very short automated sessions	Initial Access
Internal SMB activity	True Positive, High confidence	Access to an administrative share, repeated write operations, a file named after a known credential-dumping tool, large data volume and non-standard entropy	Lateral Movement

In both cases the agent recommends investigation and escalation to incident response. The SMB case is explicitly marked as not suitable for whitelisting. The web attack case adds recommendations to patch the targeted application and to block or rate-limit the source if it is confirmed malicious.

Capability 2: Deep Threat Insight

What it does. On demand, the agent correlates alert data, session metadata and real-time threat intelligence, and generates a structured brief. The vendor states this takes seconds, where manual research previously took hours.

Structure of the brief. In the vendor's remote file inclusion (RFI) example, the generated report contained seven sections plus a summary:

- | | |
|---|--|
| <ul style="list-style-type: none"> ▪ Key findings: MITRE tactic, severity (4/10), primary vulnerability, attack vector and tools used | <ul style="list-style-type: none"> ▪ Threat actors: The types of actor typically associated with the technique |
| <ul style="list-style-type: none"> ▪ Attack patterns: How the exploitation works and how scanners identify it | <ul style="list-style-type: none"> ▪ Mitigation strategies: Prevention, detection and response guidance |
| <ul style="list-style-type: none"> ▪ Indicators of compromise: A table of user agent, request pattern, HTTP method, file path and scanner test ID | <ul style="list-style-type: none"> ▪ Related tactics and techniques: A table of related MITRE ATT&CK IDs with descriptions |
| | <ul style="list-style-type: none"> ▪ Risk assessment: Threat level, impact and likelihood |

This format gives an analyst a decision-ready summary instead of a raw alert and a list of links to follow.

Business Value

Capability	Stated benefit
Network Alert Analyzer	Faster threat response, less alert noise and improved detection accuracy for SOC operations
Deep Threat Insight	Instant contextual insight, no manual lookups, faster decisions and investigation time cut from hours to seconds

Analyst Assessment

Strengths

- | | |
|--|---|
| <ul style="list-style-type: none"> ▪ Targets the right problems. Alert fatigue and missing context are among the most common complaints about NDR deployments. ▪ Framework alignment. Mapping to MITRE ATT&CK for tactics and MITRE D3FEND for response actions fits how SOC teams already work. | <ul style="list-style-type: none"> ▪ Explainable output. Every verdict comes with a rationale, key factors and a confidence level, which supports analyst trust and auditability. ▪ Closed loop with tuning. Whitelisting and stored reasoning feed back into the detection pipeline ▪ Embedded workflow. The agent works inside the existing alert view, so analysts do not switch tools |
|--|---|

Summary

Vehere's agentic layer is a practical application of AI to a clearly defined SOC bottleneck. By pairing automated alert verdicts with on-demand threat briefs, it aims to shorten the path from detection to decision while keeping analysts in the loop. The approach is credible and well aligned with industry frameworks. Its commercial impact will depend on measured evidence of accuracy and time savings, which the vendor should provide to support independent evaluation.