Pallas
AI ANALYST

ATHENA · PALLAS AI

AI Competitive Comparison

Pallas v3.0 vs Major SIEM / XDR Platforms and MSSP / MDR Providers — AI features only

July 2026 · Doc ATH-PALLAS-AI-CMP-2026-07 · Confidential — Internal Distribution Only

AI Features: Pallas vs Major SIEM / XDR Vendors

AI-capability comparison of Pallas v3.0 against the five major commercial SIEM/XDR platforms. Scope is limited to AI features specifically, not general SIEM functionality.

Positioning

PLATFORM	AI BRANDING	AI APPROACH
Pallas (Athena)	Pallas AI / Intelligence Engine	LLM-first natural-language analyst over Wazuh + Suricata + cloud sources
Splunk ES	Splunk AI Assistant (SPL)	NL-to-SPL translation + ML Toolkit anomaly detection
IBM QRadar	watsonx-powered QRadar Advisor	LLM summaries + MITRE mapping via watsonx
Exabeam	Exabeam Copilot (New-Scale)	UEBA behavioral ML + LLM threat explanations
LogRhythm (merged w/ Exabeam 2024)	LogRhythm Axon AI	Anomaly detection + assisted search, converging with Exabeam Copilot
Microsoft Sentinel	Security Copilot	GPT-4-class LLM across the Defender / Sentinel suite

Feature-by-Feature Matrix

AI CAPABILITY	PALLAS	SPLUNK	QRADAR	EXABEAM	LOGRHYTHM	MS SENTINEL
Natural-language querying	YES 3-tier router, 78 tools	YES NL-to-SPL Assistant	PARTIAL limited NL search	YES Copilot search	PARTIAL assisted search	YES Copilot NL / KQL
AI incident narratives	YES on every response	PARTIAL add-on	YES Advisor summaries	YES threat explanations	PARTIAL basic	YES strong summaries
AI-guided triage	YES triage + Jira + MTTD/MTTR	PARTIAL via SOAR add-on	YES Advisor recommendations	YES case timelines	PARTIAL basic scoring	YES guided response
Multi-source AI correlation	YES 45 strategies	YES manual SPL heavy	YES	YES behavior-centric	PARTIAL	YES Defender-centric
UEBA / ML behavioral baselining	NO rule / severity driven	PARTIAL UBA add-on	YES	YES best-in-class	YES	YES
Agentic / automated response	PARTIAL analyst-confirmed actions	YES SOAR (Phantom)	YES SOAR	PARTIAL limited	YES SmartResponse	YES playbooks + Copilot
LLM model choice / portability	YES any HF model, self-host option	NO Splunk-hosted	NO watsonx only	NO vendor-hosted	NO	NO OpenAI / Azure only
Conversation memory / follow-ups	YES per-analyst threads	PARTIAL session only	PARTIAL	PARTIAL	NO	YES
AI cost model	YES ~$12-15 / month tokens	NO premium SKU	NO premium SKU	NO bundled premium	NO	NO ~$4 / SCU-hour
Data sovereignty	YES self-hosted stack	PARTIAL cloud	PARTIAL	NO SaaS	PARTIAL	NO Azure

Where Pallas Genuinely Wins

- **NL-query depth over open telemetry —** 78 purpose-built NL tools over Wazuh, Suricata, Defender, Intune, Cloudflare, O365, AWS, GitHub and Sophos in one chat surface. Sentinel Copilot comes closest but locks into the Microsoft stack.
- **Mandatory SOC analysis —** Competitors generate summaries on demand; Pallas appends a threat assessment to every answer by design.
- **LLM independence —** DeepSeek models were swapped in production in an afternoon. Splunk / IBM / Microsoft customers cannot change their AI vendor at all.
- **Cost —** Pallas's entire monthly AI spend (~$12) is less than one hour of Security Copilot SCU billing. Total stack cost is a rounding error vs any platform here.
- **Ticket lifecycle AI integration —** Triage, Jira creation, disposition and MTTD/MTTR fields are wired end-to-end; most vendors need a separate SOAR product.

One-Line Pitch

"Pallas delivers Sentinel-Copilot-class conversational AI and automated triage on a self-hosted open-source stack, at roughly 1% of the AI cost of any commercial SIEM - trading off mature UEBA baselining and unattended SOAR playbooks, which remain roadmap items."

Pallas vs Major MSSP / MDR Providers

MDR vendors sell humans plus a platform as a service. Athena uses Pallas to deliver the same managed service - so this table is the competitive set. Key difference: MDR vendors' AI is mostly internal plumbing that accelerates their analysts; Pallas makes the customer-facing surface itself the AI.

Positioning

PROVIDER	MODEL	AI BRANDING
Pallas (Athena)	Platform powering our own managed SOC	Pallas AI
Arctic Wolf	Concierge MDR (largest by customer count)	Alpha AI
Expel	MDR, automation-heavy, transparent workbench	Expel Workbench + Ruxie bots
eSentire	MDR with threat-hunting emphasis	Atlas AI + AI Investigator
Secureworks (Sophos-owned, 2025)	Taegis MDR / XDR	Taegis AI
Rapid7	MDR on InsightIDR	AI Engine + SOC copilot features

Feature-by-Feature Matrix

AI CAPABILITY	PALLAS	ARCTIC WOLF	EXPEL	ESENTIRE	SECUREWORKS	RAPID7
NL chat with the platform	**YES** full conversational analyst	**NO** portal only	**PARTIAL** limited	**YES** AI Investigator	**PARTIAL** basic	**PARTIAL**
AI incident narratives	**YES** every response	**PARTIAL** human-written	**YES** auto-generated findings	**YES**	**YES**	**PARTIAL**
AI triage / noise reduction	**YES** + Ageleia mute rules	**YES** internal, opaque	**YES** Ruxie auto-closes ~90% noise	**YES**	**YES**	**YES**
Customer sees the AI reasoning	**YES** tool routed, confidence, counts	**NO** black box	**YES** transparent workbench	**PARTIAL**	**PARTIAL**	**PARTIAL**
AI-assisted response actions	**PARTIAL** analyst-confirmed WAF/FW/isolate	**YES**	**YES**	**YES**	**YES**	**YES**
MTTD / MTTR tracked to ticketing	**YES** Jira per incident	**YES** their portal	**YES**	**YES**	**YES**	**YES**
LLM transparency & model choice	**YES** we pick the model	**NO**	**NO**	**NO**	**NO**	**NO**
Customer data stays on premises	**YES** prompt-only egress, self-host option	**NO** full telemetry to their cloud	**NO**	**NO**	**NO**	**NO**
Pricing transparency	**YES** infra + ~$12 / month AI	**NO** per-user, opaque	**NO**	**NO**	**NO**	**PARTIAL**

Key Takeaway

The MDR vendors' AI serves their analysts - the customer sees a portal. Pallas's differentiator is that the customer-facing surface itself is the AI: the client's own team can interrogate their security data conversationally. Only Expel (workbench transparency) and eSentire (AI Investigator) are moving in that direction. For Athena as an MSSP, the pitch against Arctic Wolf or Expel is: "their AI serves their analysts; our AI serves yours too."