



SCALING RAPID7 MDR OUTCOMES WITH AI, POWERED BY AWS



CONTENTS

The Rapid7 AI Engine	3
----------------------	---

The evolving need for AI in MDR	4
---------------------------------	---

Empowering Rapid7 SOC analysts with generative AI	5
---	---


Cutting the noise with AI-powered alert triage	6
--	---

Autonomous investigation in action	7
------------------------------------	---

Enabling secure AI use for customers	8
--------------------------------------	---

Real results for MDR customers	8
--------------------------------	---

Developing on AWS systems	9
---------------------------	---



THE RAPID7 AI ENGINE

At Rapid7, we believe artificial intelligence (AI) should be transparent, human-centered, and built on deep domain expertise. The Rapid7 AI Engine is a culmination of these principles, powering scalable, adaptive, and autonomous managed detection and response (MDR) capabilities.



Our approach emphasizes:

- **Human-centricity:** AI augments, but never replaces analyst judgment.
- **Transparency:** Every AI decision should be fully explainable and auditable.
- **Expertise-driven:** Models are shaped by real-world SOC experience.

Built on secure, compliant AWS services such as Amazon Bedrock and SageMaker, the Rapid7 AI Engine is designed to bring real-world efficiency and accuracy to SecOps teams everywhere. This eBook showcases how we harness both generative and agentic AI to deliver next-generation cybersecurity outcomes.

Research, LLMs, and building with AWS Services

After extensive research, Rapid7 selected Amazon Web Services (AWS) as its partner for building AI solutions that not only power effective MDR services, but deliver the speed and scale needed to deliver cybersecurity resilience in an increasingly complex attack landscape. With the Command Platform already built on AWS, Rapid7 has the flexibility to use native services like Bedrock and SageMaker, making it easy to work with a variety of large language models (LLMs) such as Claude and BERT. This ensures the team can quickly roll out new innovations without sacrificing security and control.

The foundation: High fidelity cybersecurity data

Rapid7 has more than **20 years** of experience detecting and defending organizations against vulnerabilities and security incidents. The result is an expansive dataset – enriched by a large, global, open-source community – that enables Rapid7 to train AI models on the highest-fidelity data possible. That data is the foundation for the machine learning and AI solutions that power Rapid7 services.


THE EVOLVING NEED FOR **AI IN MDR**

Cyber attackers are leveraging AI to move faster, scale operations, and evade detection with unprecedented precision. Emerging threats like cybercrime-as-a-service – where criminals are quickly adopting AI tools to increase the volume of cyberattacks across all public and private sectors – are growing at an alarming pace. These bad actors are frequently state-sponsored and leverage AI for the purposes of automating malware, creating sophisticated deepfakes, and automating reconnaissance.

At the same time, rising alert volume and the associated analyst fatigue sink investigative speed and quality; too much time is spent

trying to fill skills gaps with a limited pool of talent; and the demand for measurable MDR outcomes is ever-expanding.

Against this backdrop, security teams and their MDR partners realize the need for a more scalable, proactive SOC – where AI agents work on behalf of analysts to perceive increasingly sophisticated threats faster and more consistently. Equally important, these advancements need to be delivered in a way that provides security operations teams with the requisite control and visibility.



EMPOWERING RAPID7 SOC ANALYSTS **WITH GENERATIVE AI**

During active investigations, every minute matters. That's why Rapid7 built an AI-powered SOC assistant using Claude LLMs through Amazon Bedrock. This initial application of AI within the Rapid7 SOC equips analysts with instant, expert-level answers in real time, surfacing relevant investigative paths, detection context, and remediation advice.

Using an intuitive chat interface, the Rapid7 MDR team leverages the SOC assistant to gain instant access to comprehensive, Rapid7 gold-standard detection and response (D&R) knowledge curated from multiple areas, ensuring investigations are concluded faster and more thoroughly than ever before.



Rapid7 also uses GenAI to:

- Draft automated MDR incident reports, saving analysts hours of post-incident documentation and ensuring consistent investigative outputs.
- Provide real-time enrichment during alert triage and investigation.

GenAI helps Rapid7 scale analyst expertise while improving decision speed, accuracy, and overall customer experience.

LLM and RAG

Retrieval-augmented generation (RAG) enhances LLMs by allowing them to access and retrieve specialist domain information from new knowledge sources that were not part of training data. This results in AI systems that can produce domain-specific, context-aware responses, without fine-tuning or training a new LLM from scratch.

CUTTING THE NOISE WITH AI-POWERED ALERT TRIAGE




Security teams today face a crushing volume of alerts. Some organizations experience up to **11,000 daily**, with **83% false positive** and **67% ignored** – equating to **3 hours per day** spent triaging alerts and not on potentially high-priority threats.¹

Our AI-powered alert triage significantly reduces the noise and alert fatigue faced by security teams, automatically dispositioning and closing benign alerts with a market-leading 99.93% accuracy. This empowers the Rapid7 SOC and SOC teams around the world to focus on higher-fidelity threats and deliver more strategic, secure outcomes.

The AI Engine leverages a sophisticated pipeline that combines domain-specific machine learning models, Bayesian inference techniques, and natural language processing (NLP). To handle more ambiguous cases, it incorporates layered GenAI triage models designed specifically to assess difficult-to-classify alerts. This evaluation ensures the AI Engine catches truly malicious activity and can proactively triage a greater volume of benign alerts without introducing false positives.

Rapid7 alert triage capabilities provide both a label and confidence in that label for new alerts, which the AI Engine uses to decide how to modify an alert – changing its priority or closing it out. Together, all of these capabilities significantly reduce analyst fatigue, sharpen alert prioritization, and shorten the time from detection to response.

¹ 2023 State of Threat Detection: The Defenders' Dilemma. Vectra AI, Inc. July, 2023.



AUTONOMOUS INVESTIGATION **IN ACTION**

At Rapid7, AI goes even further than generating outputs based on human input, or leveraging deterministic models to triage alerts. Agentic AI takes proactive, autonomous action to move MDR investigations forward. These investigative workflows are designed to operate with minimal friction and maximum transparency.

Agents handle repetitive groundwork – collecting and analyzing alert history, target history, and relevant log data – while analysts focus on complex threats and strategic decision-making. They do so in accordance with the OSCAR framework, a repeatable, rigorous approach to alert triage and resolution that ensures investigations are efficient, consistent, and complete. Here's how it works:

- **Orient:** Analysts begin by reviewing the initial alert signal and its metadata. They determine what's known, what's missing, and what context is needed to make sense of the activity.
- **Strategize:** Based on the alert type, potential severity, and initial signal clarity, the analyst forms an approach: Can the issue be resolved with current data, or will it require deeper investigation?
- **Collect:** The analyst begins data gathering. This might include historical context on the user or asset, similar past alerts, endpoint activity logs, authentication records, or threat intel correlations.
- **Analyze:** With data in hand, the analyst interprets findings: looking for behavior patterns, verifying anomaly authenticity, and cross-referencing indicators of compromise or known malicious behaviors.

- **Report:** Once the investigation is complete, the analyst summarizes their findings, outlines recommended actions, and shares outcomes with the response team or customer for final decisions and remediation.

These steps are surfaced within Rapid7's AI-powered next-gen SIEM in a way that is fully transparent and aligned with human-centered control. Agents, analysts, and customers share the same context and visibility into the logic behind AI-powered decisions, while analysts retain full oversight and authority over the decisions being made on behalf of customers



ENABLING SECURE AI USE FOR CUSTOMERS

Organizations are rapidly adopting AI tools, which means it's Rapid7's responsibility to speedily – but thoughtfully – equip security professionals with protective oversight of these applications to ensure a security-first approach to AI adoption.

Monitoring enterprise AI usage

Rapid7 provides multiple tools that enable visibility and control over AI systems. Our AI-powered next-gen SIEM, Incident Command, tracks AI service activity and is able to detect and flag ungoverned usage, commonly known as Shadow AI. Rapid7's market-leading vulnerability management solution, InsightVM, helps identify and manage the risk associated with the infrastructure used to develop or host AI models, while our cloud-native application protection platform (CNAPP) protects cloud-native pipelines involved in model training and deployment.

Governing AI development securely

To support safe and responsible AI adoption, Rapid7's Exposure Command introduces powerful capabilities for GenAI application security, enabling smarter scanning and interaction with LLMs and chatbots. This includes 25 new attacks targeting OWASP Top 10 for LLMs, enhanced R7Crawler for dynamic interface testing, and an AI-powered validation to reduce false positives. These advancements help security teams detect and fix AI-specific vulnerabilities faster.

Our tools continuously monitor for compliance drift and over-permissive access, automatically alerting security teams and initiating remediation actions where needed. This helps enforce least privilege access (LPA) and maintain a secure baseline.

REAL RESULTS FOR MDR CUSTOMERS

With Rapid7's AI Engine at work, customers experience measurable improvements in efficiency, accuracy, and operational scale. Our AI alert triage achieves an **accuracy rate of 99.93%** in identifying benign alerts, ensuring real threats stand out clearly and can be addressed faster.

99.93%

accuracy rate

200+

analyst hours reclaimed

This level of precision translates into substantial time savings – **more than 200 analyst hours reclaimed** each week within the Rapid7 SOC. These time savings empower analysts to reallocate time to more complex tasks, strategic investigations, and final decision-making – helping provide stronger protection from the biggest threats to customer environments.

For customers, these improvements come with **zero friction and complete transparency**, thanks to seamless integration and intuitive explanations of every AI action or output in the Command Platform experience.

Our commitment to innovation is also reflected in our patent activity, with **74 AI-focused patents filed to date**. Together, these outcomes showcase a SOC that moves faster, sees more, and reacts smarter; all this whether analysts are triaging alerts, investigating threats, or responding to incidents.

Rapid7's generative and agentic AI capabilities, built responsibly on AWS, help drive security operations into the future.

DEVELOPING ON AWS SYSTEMS

Rapid7 chose AWS to help build its AI Engine because it can use data that is already secured in place on AWS, rather than transferring it across boundaries to another solution. AWS offerings also enable Rapid7 to develop robust and effective solutions for generative AI security use cases. Here are the reasons why Rapid7 — and many other customers — use AWS services for GenAI solutions.

Amazon Bedrock removes the heavy lifting

With Amazon Bedrock, Rapid7 has the flexibility and builder focus it needs. It also eliminates the heavy lifting of using GenAI. Amazon Bedrock makes it easier for Rapid7 to use foundation models like Claude, Mistral, Llama, Titan, and more to build services for its MDR team and customers.

A fully managed service, Amazon Bedrock includes models for different tasks and domains that businesses can build on. Businesses get a complete platform for quickly building, refining, and deploying GenAI applications. Amazon Bedrock also helps ensure that customer data stays under customer control. When a foundation model is tuned, it's based on a private copy of that model. Data is not shared with model providers and is not used to improve the base models. Amazon Bedrock offers comprehensive monitoring and logging capabilities that can support governance and audit requirements.

What are Foundation Models?

Foundation models are large AI models trained on vast amounts of data. These models serve as a base that can be adapted and fine-tuned for various generative AI tasks and applications. By providing a strong starting point, foundation models accelerate the development of AI systems for different industries and use cases.



Amazon SageMaker accelerates development

With Amazon SageMaker, Rapid7 can accelerate the deployment and fine-tuning of its models. The platform also provides pre-trained foundation models and tools to developers so they can quickly build GenAI applications. The service makes it easy to experiment with prompts, adjust model settings, and deploy a GenAI app with just a few clicks.

AWS AI services are secure and responsible by design

AWS enables enterprises to scale, migrate, and manage AI applications, models, and workloads on a global cloud infrastructure that is secure by design. It supports **143** security standards and certifications — including PCI-DSS, HIPAA/HITECH, FedRAMP, GDPR, FIPS 140-2, and NIST 800-171. Also, AWS is committed to responsible AI through its work with academic, industry, and government organizations.



LEARN MORE

Contact Rapid7 for more information about the AI-powered Command Platform.