EBOOK

Rapid7 Supercharges SecOps with Generative AI Solutions Built on AWS

Table of Contents

Building Generative AI Solutions to Enhance Cybersecurity

Artificial Intelligence (AI)—and specifically generative AI—is having a profound impact on every industry, and cybersecurity is no exception. Generative AI analyzes complex and unstructured historical data, using context and insights to recognize new and evolving threats in real time. These techniques ensure defenders are more efficient and effective at distilling disparate sources of complex information into specific, actionable guidance to prevent and remediate security incidents.

With the emergence of generative AI, Rapid7 embraced the opportunity to use this technology, along with existing machine learning (ML) capabilities, to enhance the Insight Platform and services it provides to managed services customers.

Research, LLMs, and Building with AWS Services

After extensive research, Rapid7 selected Amazon Web Services (AWS) for building generative AI solutions. With the Insight Platform already built on AWS, Rapid7 has the flexibility to use services like Amazon Bedrock and Amazon SageMaker. These services make it easier to work with a variety of large language models (LLMs), such as Claude and BERT. As a result, the Rapid7 team can quickly roll out new innovations without sacrificing security and control.

Generative AI that Powers Cybersecurity and Secure AI Application Development

Rapid7 uses Amazon Bedrock with Claude to develop key parts of the Rapid7 generative AI engine, which helps security teams improve threat detection efficacy and operational efficiency. This engine supercharges security operations, allowing the Rapid7 Managed Detection and Response (MDR) team to deliver better security outcomes on the frontline. Furthermore, it provides added security context in the moment during an investigation and generates incident reports for the Rapid7 Security Operations Center (SOC).

These capabilities are only part of the story. This ebook explains how Rapid7 uses generative and more conventional AI to improve the cybersecurity services it provides to MDR customers. It also explains how Rapid7 solutions help customers securely develop their own generative AI models and seamlessly integrate with their infrastructure.

The Foundation: High-Fidelity Cybersecurity Data

Rapid7 has more than 20 years of experience defending organizations against vulnerabilities and security incidents, creating multiple high-fidelity datasets. Enriched by the global open-source community, this expansive data is the foundation for Rapid7's generative AI and AI/ML solutions.

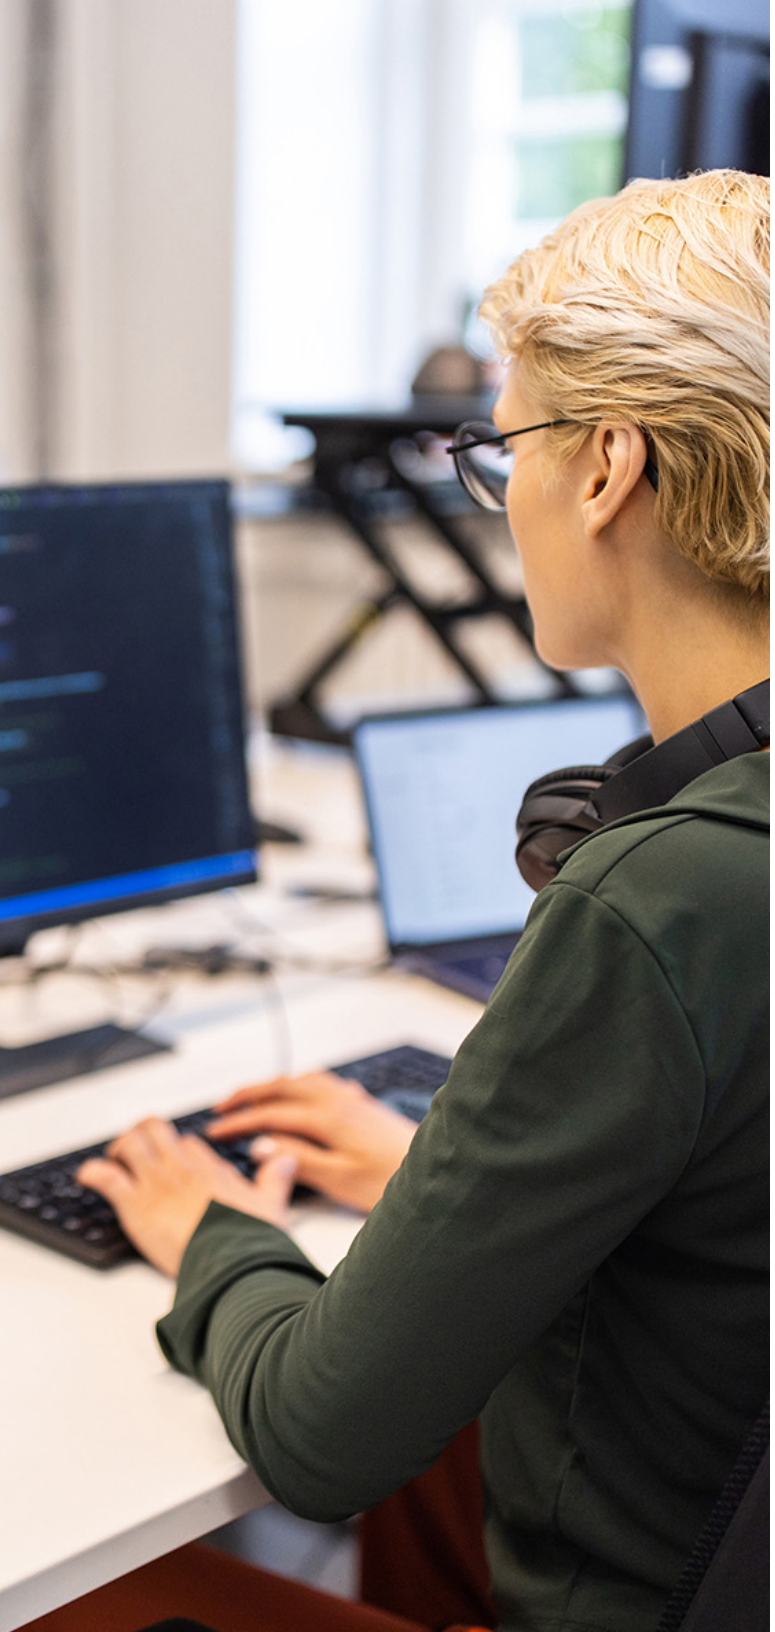

Improving Security Management in the Rapid7 SOC with AI

Rapid7 infuses AI into its SOC workflows to improve how MDR analysts detect security threats and proactively mitigate risk. These features, including those built with generative AI, are regularly informed by a continuous feedback loop with Rapid7 analysts and their customers.

The Generative AI-Powered SOC Assistant

Time is of the essence during an investigation. Rapid7's AI assistant ensures SOC teams can keep pace by quickly providing expert guidance on a wide range of complex detection and response (D&R) situations.

Built by the Rapid7 in-house AI team using Amazon Bedrock, Claude LLMs, and retrieval-augmented generation (RAG), this SOC assistant enhances the efficiency of analyst investigations. Using an intuitive chat interface, the Rapid7 MDR team has instant access to comprehensive, Rapid7 gold-standard D&R knowledge curated from multiple areas. This ensures that investigations are concluded faster and more thoroughly than ever before.

Large Language Models (LLMs) and Retrieval-Augmented Generation (RAG)

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 from scratch.

Automated MDR Report Drafting with Generative AI

Rapid7 is using generative AI to carefully automate the drafting of security reports for SOC analysts, normally a completely manual process. With more than 11,000 customers globally, the Rapid7 SOC triages a huge volume of activity from month to month, with timely summaries that are critical for keeping customers fully updated. The Rapid7 AI team conducts meticulous research to fine-tune LLMs that can automatically summarize D&R incidents. This valuable feature streamlines the writing process, saving significant time to help produce MDR reports at scale. A Rapid7 expert reviews the reports and validates them for accuracy. Thanks to this accelerated report delivery to customers, analysts can get back to the frontline sooner.

Improve Detection Efficacy and Eliminate False Positives with AI-Assisted Alert Triage

The Rapid7 AI-based detection engine provides an alert triage function. Using historical security data qualified by their expert security analysts, and both traditional ML and generative AI models, Rapid7 AI-assisted alert triage can accurately label new security alerts as malicious or benign. Such labeling upgrades the priority and visibility of alerts qualified as malicious, and proactively triages alerts qualified as benign. This work boosts the signal and reduces the noise, thereby enabling Rapid7 analysts to spend more time investigating the security signals that matter to their customers.

The Foundation of AI-Assisted Alert Triage

For alert triage, the Rapid7 AI engine employs domain-specific text processing, Bayesian statistics, and techniques borrowed from natural language processing to generate features used to build traditional ML models. These models provide both a label and confidence in that label for new alerts. The Rapid7 AI engine uses them to decide how to modify an alert by changing its priority or closing it altogether.

Multilayered AI Alert Triage with Generative AI

To reduce the likelihood of misclassification, the primary AI engine alert triage function is conservative in its interpretation of confidence scores. This can leave some alerts unprocessed, requiring manual intervention. Recently, Rapid7 developed a specialized secondary set of AI models, purpose-built to process alerts with more variability. Using fine-tuned language models based on BERT and built for classification, the alert triage now routes difficult-to-classify alerts to a specialized layer to perform a secondary evaluation of unprocessed 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.

Enabling Customers to Manage and Monitor Enterprise Generative AI Services

Many companies are using, building, and integrating generative AI for internal and commercial purposes. These services pose new risks caused by a lack of oversight, incomplete usage policies, and insufficient monitoring. Rapid7's AI best practices are encapsulated across InsightIDR, InsightVM, and InsightCloudSec to help enterprises strengthen their security postures, including the development and consumption of generative AI.

InsightIDR has services for monitoring and protecting on-premises and cloud resources, including configurable alerting for targeted AI SaaS vendors. This gives Rapid7 the ability to detect the ungoverned use of consumer AI, a footprint known as shadow AI, which is a fast-emerging threat in enterprise environments.

InsightVM monitors systems used to develop or host AI tools, as well as endpoints and development labs used for research. InsightCloudSec provides mechanisms to monitor and protect cloud-based systems used to develop generative AI, plus it provides visibility and effective guardrails based on best practices.

Secure Development of Generative AI Applications in the Cloud

InsightCloudSec supports secure development of generative AI in the cloud with real-time and continuous visibility into AI/ML resources running across an AWS environment.

Organizational policies informed by Rapid7's AI/ML security best practices help teams understand what controls to implement and provide the ability to enforce compliance as developers begin building generative AI applications. Aligned to the OWASP Top 10 for ML and LLMs, these best practices continuously audit a comprehensive inventory of all assets across container and cloud environments in real time to consolidate policy management and eliminate noncompliant configurations. The platform goes beyond simply providing visibility into the resources themselves. It also includes which users across the organization have access to them, with the ability to right-size permissions in accordance with least privilege access (LPA).

Findings are enriched by Layered Context. This provides security teams with a unified, contextualized view of their AI-related risk in a single place, which enables companies to innovate and accelerate at scale. With the added context, organizations are equipped with the information needed to effectively prioritize signals based on exploitability and potential business impact. Native automation alerts both security teams and developers to compliance drift and even offers the ability to auto-remediate risky configurations or overly permissive access risk as soon as it's detected.

Rapid7 Leverages InsightCloudSec to Securely Build Generative AI Solutions

To ensure the security and compliance of its generative AI solutions, Rapid7 uses its proprietary platform to implement and enforce policies and maintain best practices. Rapid7's models are also built, tested, and continuously trained by its internal SOC experts. Rapid7 builds, tests, and continuously augments its models with data it collects, solutions it offers customers, and internal MDR and SOC. For example, to ensure the efficacy of the insights derived by the generative AI engine, it was trained by Rapid7's world-renowned AI security practitioners.

Visibility and Effective Guardrails for Safely Building and Consuming Generative AI

With services that span content moderation and translation to model customization, the Rapid7 AI/ML Security Best Practices compliance pack continuously assesses customer environments. Through event-driven harvesting, the compliance pack ensures developer teams have the appropriate guardrails to develop with generative AI. Among its diverse array of services are vulnerability controls and continuous permission and usage monitoring.

Vulnerability Controls

The Rapid7 AI/ML Security Best Practices compliance pack introduces controls for the 11 vulnerabilities in this table. With this pack, it's possible to check alignment with each of these controls in one place.

Data poisoning Manipulates training data to influence how a model makes decisions	**Model poisoning** Manipulates the model itself to influence how it makes decisions	**Model stealing** Gains access to the model and its parameters
Model inversion Reverse-engineers a model to understand how it can be influenced	**Transfer learning** Trains a model on a specific task, but fine-tunes it on another to influence its decisions	**Model skewing** Manipulates training data distribution to influence a model's decisions
Excessive permissions Exploits overly permissive roles and policies to manipulate systems, services, and models	**Denial of service** Drives service degradation and increased cost	**Supply chain compromise** Modifies or replaces a third-party library used by a system or its data
Membership inference Manipulates training data to reveal sensitive information such as personal identifiable information and protected health information	**Output integrity** Modifies the output from an ML model to negatively impact business processes or systems	

Continuous Monitoring for a Holistic View

Continuous monitoring for effective permissions and actual usage helps right-size permissions to ensure alignment with LPA. Combined with other components, the result is a holistic view of a customer's compliance landscape that facilitates better strategic planning and decision-making. Customers can also set up automated alerting and remediation for drift detection and prevention mechanisms.

As mentioned previously, Rapid7's solutions and services for secure generative AI development and guardrails are built with AWS services. They also run on AWS infrastructure that is secure by design.

Why AWS for Generative AI

Rapid7 chose AWS to help build its generative 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 generative AI solutions.

Amazon Bedrock Removes the Heavy Lifting

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

A fully managed service, Amazon Bedrock includes models for different tasks and domains that businesses can build on. With Amazon Bedrock, businesses get a complete platform for quickly building, refining, and deploying generative AI 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.

Amazon SageMaker Accelerates Development

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

AWS Generative AI Services are Secure and Responsible by Design

AWS enables enterprises to scale, migrate, and manage generative 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.

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.

Learn More

Contact us for more information about the Rapid7 AI-powered SecOps platform.

© RAPID7 2024