

How Tenable Uses AI and Machine Learning

Last Revised: October 20, 2026



Table of Contents

Introduction	4
At a Glance: Where Your Data Meets AI	5
Current Use of AI for Exposure Management	7
Machine Learning	7
Deep Learning and Small Language Models	7
Tenable One AI Exposure	7
Generative and Agentic AI	8
Prioritizing Vulnerabilities	8
Assessing Attack Path Risks	8
Managing Inventory	9
Monitoring Exposure Signals	9
Utilizing Tenable Hexa AI	9
How Tenable Research Uses AI	10
Researcher-Augmented Content Writing	11
Agentic Content Writers	11
Multi-Language Detection Evaluation	12
Risk Classification of AI Agent Tools	12
Development Pipeline Workflow	13
Training Data	14



Third-Party AI Providers	15
Data Retention and Governance	15
Examples of Features and Technical Specifications	16
Conclusion	17
Relevant Links	18



Introduction

Tenable uses machine learning and generative AI to strengthen its cybersecurity products in three areas: (1) proactively identifying and prioritizing cyber exposure risk; (2) increasing the efficiency of security operations; and (3) giving customers clear visibility and insight to prioritize and reduce risk.

- [At a Glance: Where Your Data Meets AI](#)
- [Current Use of AI for Exposure Management](#)
- [How Tenable Research Uses AI](#)
- [Training Data](#)
- [Third-Party AI Providers](#)
- [Examples of Features and Technical Specifications](#)
- [Conclusion](#)
- [Relevant Links](#)



At a Glance: Where Your Data Meets AI

This table summarizes every Tenable AI-powered feature described herein.

Feature	AI/ML Technology	Customer data usage	Third-party LLM usage
Vulnerability Priority Rating (VPR)	Machine Learning / Generative AI	No: scores are computed from public and commercial threat intelligence, not your environment	Yes, for the generative-AI component that reads news and advisory content (not your data)
CVSS Metrics Prediction	Machine Learning	No	No
Operating System Prediction	Machine Learning	Yes: your scan data (network-level host metadata only)	No
User Criticality Rating	Machine Learning	Yes: your users' job titles and access rights	No
AI Threat and Misuse Detection	Deep Learning / Small Language Models	Yes: activity on your connected AI platforms	No (Tenable-trained models)
Data Exposure Detection	Deep Learning / Small Language Models	Yes: content passing through your connected AI platforms	No (Tenable-trained models)
Prompt Translation	Deep Learning / Small Language	Yes: prompts you submit in a non-English language	No (Tenable-trained models)



Feature	AI/ML Technology	Customer data usage	Third-party LLM usage
	Models		
Attack Path Summary, Assistant, and Mitigation Guidance	Generative AI	Yes: your attack path, asset, and vulnerability data	Yes (Google Vertex AI)
Asset Summary and Natural Language Search	Generative AI	Yes: your asset inventory	Yes (Google Vertex AI)
Exposure Signals Description and Summary	Generative AI	Yes: your exposure signal data	Yes (Google Vertex AI)
Tenable Hexa AI	Generative / Agentic AI	Yes: your asset, vulnerability, identity, configuration, and AI-system data already in Tenable, plus what you type	Yes (Anthropic, via AWS Bedrock)
Adversary View (Tenable Hexa AI)	Generative / Agentic AI	Yes: your exposure, asset, and vulnerability data already collected in Tenable One	Yes (Anthropic Claude Mythos, via AWS Bedrock)



Feature	AI/ML Technology	Customer data usage	Third-party LLM usage
Tenable Research	Generative / Agentic AI	Yes: your scan data (network-level host metadata only)	Yes, (Anthropic, via Claude Code)

Current Use of AI for Exposure Management

The sections below describe each technique in more detail.

Machine Learning

Tenable uses machine learning to proactively identify and prioritize vulnerabilities, helping you focus on the most critical threats:

- **Vulnerability Priority Rating (VPR):** Predicts the likelihood that a vulnerability will be exploited within 28 days, enabling risk-based prioritization. Available across the Tenable One platform.
- **CVSS Metrics Prediction:** Predicts CVSSv3 metrics, providing a key input to the VPR model.
- **Operating System Prediction:** Improves remote operating system fingerprinting, even with limited scan data, leading to better device management. Available as Tenable Nessus plugin 132935.
- **User Criticality Rating:** Assesses the importance of individual users within your organization based on their job titles and access rights, enabling risk-based exposure management within Tenable One Identity Exposure and Tenable One.

Deep Learning and Small Language Models

Tenable One AI Exposure



Tenable One AI Exposure uses deep learning models and fine-tuned small language models (SLMs) to secure interactions across any AI platform, including chatbots, agents, and applications.

- **AI Threat and Misuse Detection:** Identifies potential adversarial attacks targeting your AI platforms, and detects misuse such as unauthorized data access or unsafe reliance on AI for decisions.
- **Data Exposure Detection:** Detects sensitive data that an AI engine or agent exposes or leaks.
- **Prompt Translation:** Uses AI to translate prompts from non-English languages to English for analysis.

These models are trained by Tenable using open-source libraries. Your AI activity is evaluated at runtime against pre-built detection policies to identify risk; that activity is not used to train or modify the models.

Generative and Agentic AI

Tenable leverages generative AI and pre-trained Large Language Models (LLMs) from Google's Vertex AI and Amazon's Bedrock platforms to deliver clear, concise, and actionable information, including for:

Prioritizing Vulnerabilities

- **Vulnerability Priority Rating (VPR):** Analyzes vast amounts of content from cyber news media, security advisories, vendor blogs, and related sources to help predict the likelihood of exploitation within 28 days. Incorporates context-rich vulnerability insights as inputs to the threat model, updates key drivers, and enriches metadata, including threat and remediation summaries. Available across Tenable One.

Assessing Attack Path Risks



- **Attack Path Summary:** Distills complex attack paths into easy-to-understand summaries.
- **Attack Path Assistant:** Lets you ask natural language questions about specific attack paths and their components.
- **Attack Path Mitigation Guidance:** Offers step-by-step instructions for addressing identified attack path risks.

These features summarize and explain attack paths that Tenable's existing Attack Path Analysis has already identified. They are distinct from Adversary View, a Hexa AI capability described below that uses Claude Mythos to reason across raw exposure data and surface new attack chains, rather than explain ones already found.

Managing Inventory

- **Asset Summary:** Provides detailed descriptions of your assets and highlights their most critical exposures.
- **Natural Language Search:** Lets you search your asset inventory using natural language.

Monitoring Exposure Signals

- **Exposure Signals Description:** Explains specific exposure insights.
- **Exposure Signals Summary:** Explains why an asset is included in a particular exposure insight.

Utilizing Tenable Hexa AI

Tenable Hexa AI uses agentic AI to plan and carry out multi-step security workflows over your own Tenable data. It breaks complex requests into a sequence of steps and takes action through Tenable's existing platform capabilities, working against the vulnerability, asset, identity, configuration, and AI-system data your Tenable products have already collected, plus whatever you type into it directly.



Tenable partners with Anthropic to deliver these agentic capabilities across Tenable One. No customer data is used to train the underlying models.

Users interact with Tenable Hexa AI through two channels:

- **Chat:** A conversational interface embedded in Tenable One, Tenable One Vulnerability Management, and Tenable One Cloud Exposure. Users describe what they want to accomplish in natural language; Hexa AI plans the steps, surfaces its reasoning, and asks for confirmation before making any change.
- **Model Context Protocol (MCP):** A Tenable-hosted endpoint that exposes Hexa AI capabilities to MCP-compatible AI assistants of the customer's choice, letting customers use Hexa AI through the AI client they already use.

Hexa AI capabilities drive common security workflows and Tenable product features, including:

- **Asset Tagging and Organization:** Search the asset inventory using natural language, create tags and tag categories, and apply tags to matched assets at scale.
- **Exposure Investigation and Organization:** Search and inspect vulnerabilities and exposures across Tenable One; review affected assets and asset details or adjust finding risk.
- **Dashboard and Report Creation:** Build, update, or remove dashboards on demand.
- **Scan Operations:** Create, configure, launch, and monitor scans and scan policies.
- **Remediation Workflows:** Create, update, or remove Combinations, Initiatives, and leverage third-party integrations to create ticketing or change-management artifacts.
- **Adversary View (Attack Path Discovery):** Powered by Anthropic's Claude Mythos model, reasons across the exposure, asset, and vulnerability data already collected in Tenable One to reconstruct how an attacker could move from an initial point of access to critical systems, surfacing the evidence behind each step and recommending the remediation action that would break the path. Unlike the tagging, dashboard, and remediation-workflow capabilities above, Adversary View is read-only; it analyzes and recommends, but does not modify any data itself.

How Tenable Research Uses AI



Tenable Research has integrated AI-assisted workflows to address the accelerating pace of vulnerability disclosures and the demand for high-quality exposure intelligence, spanning all sensors across Tenable One. This work aims to provide value in three areas: faster time to quality coverage of exposures and contextual intelligence, new features that help customers prioritize and remediate, and frontier research that surfaces emerging threats, attack surfaces, and technologies that will shape the future of exposure management.

Tenable Research (the team that builds the detections and vulnerability content behind Tenable's plugin feed) uses AI throughout its content development process. This is a researcher-facing use: it runs inside Tenable's development environment to help build detection content, and the AI models involved are not packaged into or distributed as part of the plugin feed, audit files, or other content customers receive.

Researcher-Augmented Content Writing

Individual researchers use AI tooling locally to speed up manual content development, automated feed development, and exposure intelligence collection and analysis, primarily for more complex development tasks. Tenable Research continues to develop and share new AI skills, agents, and tools across the team to take on more of this work.

Agentic Content Writers

Tenable Research has integrated agentic content-writing skills into its development pipeline. In this model, an AI agent handles the bulk of initial content development, and a researcher is responsible for review, approval, and release. This augments the manually written content track rather than replacing researcher judgment on quality and accuracy.

Translating a vendor advisory into executable detection content is not always straightforward: vendor advisories often contain caveats and workarounds that require human interpretation. Tenable Research uses AI agents and agentic judge models to reach high-accuracy consensus on a vendor's intent, producing automated plugin coverage that would not otherwise be attainable.



A small percentage of vulnerability configuration coverage is released without human review, but only when both of the following are true:

- The content matches a strict allowlist of approved vendor configuration commands or regular expressions, which prevents the generation of any malicious commands. All commands are strictly read-only "show" commands that cannot change any configuration.
- Multiple independent agentic judge models agree on the result with very high confidence.

If either condition is not met, the content is routed to a human reviewer before release.

Multi-Language Detection Evaluation

Tenable Research uses AI to extend its AI Exposure detection models to additional languages, running automated translation pipelines across its evaluation and training sets. This allows Tenable to build detections with high fidelity across multiple languages and a wider range of security use cases.

Risk Classification of AI Agent Tools

As organizations connect AI agents to enterprise data through platforms such as Microsoft 365 Copilot Studio, Tenable Research uses AI to assess the security risk each available agent tool introduces. A single AI agent can invoke hundreds of connectors (for example, Salesforce, OneDrive, or Dataverse), and each connector exposes many individual actions, such as updating a row, retrieving a file, or deleting an item. Determining which actions are risky in the hands of an autonomous agent is a high-volume, expert-judgment task well suited to AI-assisted review.

Several guardrails preserve researcher oversight:

- **Governed question lifecycle:** New risk questions start in an evaluation state and only affect live risk scoring after a researcher promotes them to production.



- **Evaluation against a human-labeled golden dataset:** Before promotion, the pipeline re-runs classification against a curated set of actions whose correct risk status was set by human experts, and reports the precision impact of the proposed change and exactly which actions it would affect.
- **Human review and override:** Researchers can inspect every action, its AI-generated summary, and its classification, and can override the automated result: a human decision always takes precedence over the model, so an action a researcher marks as not risky is excluded regardless of its computed score, and vice versa.
- **Coverage and redundancy analysis:** Tooling reports how many risky actions each question catches, helping researchers keep the question set high-signal and free of redundancy.

Development Pipeline Workflow

All Tenable Research content passes through a CI/CD pipeline that automatically tests the quality of every new or modified plugin, plus a peer review process in which multiple reviewers must approve the underlying NASL code before it merges into the feed and ships in Tenable products. The one exception is the narrow agentic-judge path for vulnerability configuration coverage described above, under Agentic Content Writers. As with the rest of this document, AI models used in this pipeline are not packaged into or distributed as part of any content that reaches customer environments.

For more on Tenable's plugin QA/QC process, see [Customer Outreach on Tenable Plugin QA/QC Processes](#).



Training Data

Tenable leverages various data sources, including public, commercial, and anonymized internal information, to train its machine learning models. Tenable does not use identifiable customer data to train any model, including large language models, which is designed to eliminate the risk of data leakage.



Third-Party AI Providers

In certain cases, Tenable purchases AI services to train Tenable-specific models or build Tenable AI solutions. Tenable only uses generative AI in its applications through a platform approved by Tenable's internal AI Governance Council, which regularly reviews Tenable's corporate and product use of AI technologies and methods.

Google Cloud's Vertex AI platform powers the generative AI features described above for Tenable One (see [Current Use of AI for Exposure Management](#)). Amazon Web Services (AWS) Bedrock and Anthropic power Tenable Hexa AI, providing foundation models for agentic reasoning, planning, and tool use, along with content safety and policy enforcement. Adversary View, specifically, uses Anthropic's Claude Mythos model family for its attack-path reasoning.

Tenable's approved third-party providers access customer data solely to generate a response, under contract terms that prohibit them from retaining that data or using it to train their own models. Data submitted for inference is not logged, cached, or reused by the provider beyond generating that response. The one exception is Adversary View (powered by Anthropic's Claude Mythos model): Anthropic retains submitted data for 30 days for trust-and-safety purposes, rather than processing it strictly at inference time. Even during this retention window, Anthropic does not use the data to train its models.

Data Retention and Governance

Customer data and inputs or prompts processed in connection with ML or generative AI product features are retained only as long as reasonably required to provide the requested services. The capabilities described above are applied to customer-collected data that resides within the Tenable platform, which is governed by the [Tenable Master Services Agreement](#). More information about how Tenable protects customer data is available at tenable.com/trust/assurance.



Examples of Features and Technical Specifications

One example of Tenable's use of AI is the Vulnerability Priority Rating (VPR) model, which uses machine learning to predict the likelihood of a CVE being targeted for exploitation in the next 28 days. This dynamic model runs every 24 hours, taking into account the latest threat and vulnerability data for updated VPR scores. The model is trained using a Random Forest classifier that uses CVEs exploited within a 28-day period as the target variable.

Another example is the Attack Path Summary feature, which utilizes a pre-trained large language model to provide a contextual summary of specific attack paths. Users can provide feedback using a thumbs-up/thumbs-down system to improve the accuracy and effectiveness of the feature. For more information, see the [relevant links](#) below.

Further, Tenable Hexa AI is an example of applying agentic AI. When a user asks Hexa AI to "tag all my Windows Server assets and build a compliance dashboard," Hexa AI breaks the request into a sequence of steps: search the asset inventory for Windows Server assets, identify or create the appropriate tag, request user approval to apply the tag at scale, apply the tag upon confirmation, and create a dashboard from a relevant template with the new tag applied as a context filter. Each step that changes data (applying tags, creating a dashboard) requires explicit user approval before it runs. Hexa AI operates against the same Tenable platform capabilities a user would access directly, so all actions respect the user's existing access permissions and are fully auditable.



Conclusion

Tenable has integrated a range of AI techniques into its product portfolio and its research operations: traditional classification models, generative AI, and agentic AI that carries out multi-step workflows with human oversight, alongside an AI-assisted content development pipeline that keeps a human reviewer in the loop for the plugin feed customers receive. Tenable's AI Governance Council reviews these uses on an ongoing basis.



Relevant Links

- [Threat score prediction model \(Google Patents\)](#)
- [Automatic generation of vulnerability metrics using machine learning \(Google Patents\)](#)
- [Host operating system identification using transport layer probe metadata and machine learning \(Google Patents\)](#)
- [Predictive Prioritization: How to Focus on the Vulnerabilities That Matter Most \(Tenable whitepaper\)](#)
- [What Is VPR and How Is It Different from CVSS? \(Tenable blog\)](#)
- [How VPR Helped Prioritize the Most Dangerous CVEs in 2019 \(Tenable blog\)](#)
- [How to Use VPR to Manage Threats Prior to NVD Publication \(Tenable blog\)](#)
- [Not a Blackbelt in Attack Path Analysis? Tenable ExposureAI Helps You Achieve Proactive Security \(Tenable blog\)](#)
- [Crushing the Axios supply chain threat with Tenable Hexa AI: Use cases for agentic AI](#)
- [Beating the Mythos clock: Using Tenable Hexa AI custom agents for automated patching](#)
- [Vulnerability remediation: Match CVEs to asset owners in seconds with Tenable Hexa AI](#)
- [Customer Outreach on Tenable Plugin QA/QC Processes](#)
- [Tenable to Bring Claude Mythos 5 into the Tenable One Exposure Management Platform](#)