




# Tenable Cloud AI Risk Report 2025

Overlooked misconfigurations, risky defaults in managed services and insights for secure AI adoption





# Table of contents

<b>Introduction</b>	<b>03</b>
<b>Executive summary</b>	<b>04</b>
<b>Key findings</b>	<b>05</b>
<b>AI adoption</b>	<b>06</b>
<b>Adoption of AI developer packages</b>	<b>07</b>
<b>Adoption of managed cloud AI developer services</b>	<b>08</b>
<b>AI risks and misconfigurations in the cloud</b>	<b>09</b>
→ <b>Unremediated critical vulnerabilities</b>	<b>10</b>
→ <b>Jenga concept meets AI</b>	<b>11</b>
→ <b>Amazon Bedrock training bucket without public access blocked</b>	<b>12</b>
→ <b>Amazon Bedrock training buckets are overly permissive</b>	<b>13</b>
→ <b>Amazon SageMaker with root access enabled</b>	<b>14</b>
<b>Mitigation strategies for AI risks</b>	<b>15</b>
<b>Closing thoughts</b>	<b>16</b>
<b>Methodology</b>	<b>16</b>
<b>About Tenable Cloud Research</b>	<b>16</b>

# Introduction

Artificial intelligence (AI) is here, helping organizations improve their efficiency, decision-making and competitive advantage. The gain comes with new security challenges. AI tools can propagate security flaws; sensitive AI assets, deeply integrated within business operations, can contain vulnerabilities or misconfigurations that pose risks. As part of a mature exposure management strategy, security stakeholders must understand these AI risks and take proactive steps to not only secure their AI tools and resources but also prevent them from creating risky exposures in their cloud environment.

This report draws on Tenable Cloud Research's analysis of workloads and assets across diverse cloud and enterprise environments to highlight the current state of security risks in cloud AI development tools and frameworks, and in AI services offered by the three major cloud providers — Amazon Web Services (AWS), Google Cloud Platform (GCP) and Microsoft Azure. We provide guidance for raising risk awareness among security and developer teams, identifying blindspots and otherwise protecting your cloud environment as you adopt AI technologies.

# Executive summary

As organizations expand their use of AI, decision makers need to adapt their cybersecurity and compliance strategies to include AI-related tools and data. This includes ensuring that teams and individuals from software engineers to DevOps and security teams understand AI's unique (and not so unique) security pitfalls. The Tenable Cloud AI Risk Report 2025 highlights key, often overlooked, AI security risks that affect AI services, software tools and applications, as well as AI training data — and offers actionable insights for their mitigation.

One serious risky pattern we found is the occurrence of the Jenga® concept in managed AI services. The Jenga concept identifies the tendency of cloud providers to build one service on top of the other, with “behind the scenes” building blocks inheriting risky defaults from one layer to the next. Such cloud misconfigurations, especially in AI environments, can have severe risk implications if exploited.

Another finding that gives cause for concern is the common misconfiguration of overprivileged identities in AI implementations — and an even greater risk when in toxic combination with a critical vulnerability or public exposure. Some AI services offer fine-grained access control. Such permissions management capabilities, though, are per cloud service and cloud provider only. Configuring sufficient permissions while adhering to least privilege requires contextual insight across cloud (and multi-cloud) environments, including workloads, identities, sensitive data and other resources.

In this report, we present the AI risks we have observed in self-managed AI developer tools and AI cloud services, and provide mitigation and best practice recommendations for securing AI in the cloud. To meet the new AI challenges, it is essential that organizations take a cloud-native application protection platform (CNAPP) approach that contextualizes and prioritizes risk across cloud infrastructure, workloads, network, identities, data and AI. To cut to the chase: AI, for all its intelligence, is not risk-free and requires your attention.



# Key findings

Our analysis of AI in cloud environments revealed adoption levels and risky patterns in select tools and services.



## 70% of cloud AI workloads contain unremediated critical vulnerabilities

Nearly three-quarters of cloud workloads with AI software installed – across Azure, AWS and GCP – have at least one critical vulnerability (higher than non-AI workloads), making them prime targets for attackers.



## AI adoption is under way, introducing new challenges in securing cloud workloads

Many are turning to cloud-based AI services to power their business; 60% of Azure users have configured its Cognitive Services, 25% of AWS users have configured Amazon SageMaker and 20% of GCP users have configured Vertex AI Workbench in Google Cloud. AI technologies increase cloud data volume and sensitivity, which can raise security and compliance risk.



## Jenga-style cloud misconfigurations are surfacing in AI services

Cloud providers are layering AI services on top of each other, creating building blocks often unknown to users that inherit hard-to-detect-and-fix risky default settings. 77% of the organizations that have Vertex AI Workbench in Google Cloud set up have at least one notebook instance configured with the overprivileged default Compute Engine service account.



## Risky default permissions are introducing unnecessary AI risk

The vast majority, 91%, of the organizations with Amazon SageMaker set up have the risky default of root access (i.e. administrator privileges) in at least one notebook instance – enabling users to change system-critical files including those contributing to the AI model.



## Excessive exposure is putting AI training data at risk

14% of the organizations with Amazon Bedrock set up have at least one AI training bucket configured to not prevent public access; 5% have at least one overly-permissive AI training bucket. This potentially exposes AI training data to risks such as poisoning, model manipulation and sensitive data leakage.

# AI adoption

Organizations are installing AI developer packages, such as Scikit-learn and TensorFlow, to build and automate AI applications on their own workloads. Cloud AI services configured by organizations include Azure Cognitive Services, Azure AI Bot Service, Amazon SageMaker and Vertex AI Workbench in Google Cloud. Such services help organizations simplify AI adoption efforts through automation, ease of use and scalability so they can tailor pre-existing AI models and workflows to their needs and applications.

After years of little change, AI adoption had a dramatic spike last year.



## AI adoption worldwide has increased dramatically in the past year, after years of little meaningful change

Organizations that have adopted AI in at least 1 business function, % of respondents




Figure 1 - McKinsey reported a dramatic spike in AI adoption by organizations in early 2024, in contrast with the seven years prior. The 2024 data, collected between February 22 and March 5, 2024, involved 1,363 participants across regions, industries, company size, functional specialties and tenures. (Source: [McKinsey Global Survey on AI](#))

# Adoption of AI developer packages

AI developer packages (also called frameworks) address different needs, including machine learning, Natural Language Processing (NLP), image processing and computer vision.

## Deployment rates of AI packages — by package name

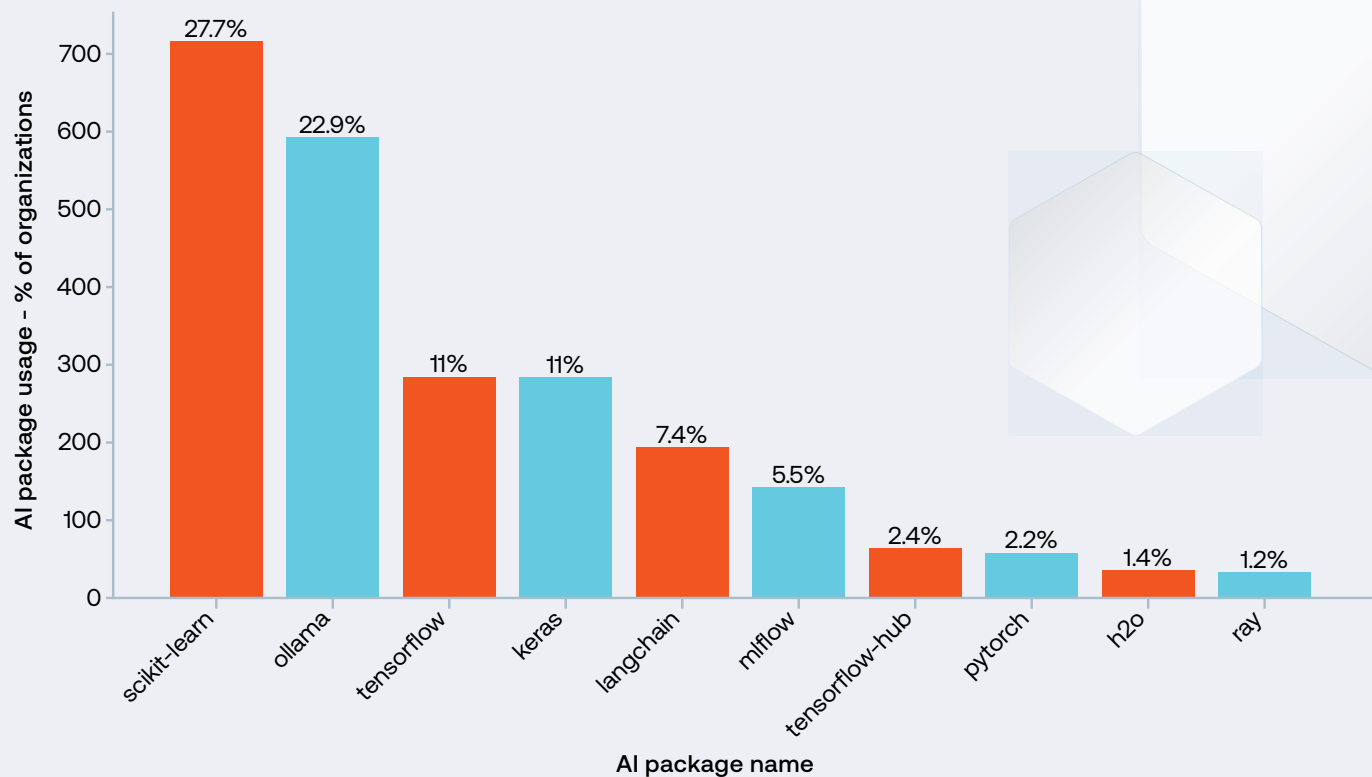


Figure 2 - Scikit-learn (27.7%) and Ollama (22.9%) are the most widely deployed self-managed AI development frameworks for machine learning among the organizations we studied

Among the organizations that have AI tools installed, we found Scikit-learn, Ollama, TensorFlow and Keras to be the tools most widely deployed for machine learning. Organizations deploy Scikit-learn for traditional machine learning, Ollama for integrating language models into applications and TensorFlow and Keras for deep learning.

Like most of the other self-managed AI packages observed, Scikit-learn and Ollama are open source. Scikit-learn, released

in 2010, is a foundational, Python-based machine learning tool. Ollama, released in July 2023, allows running and having AI models interact locally, regardless of specific frameworks. It is likely that organizations are deploying these two tools more than the other tools due to Scikit-learn's long-time presence and Ollama's enabling of work with AI models in the privacy of a non-cloud environment.

# Adoption of managed cloud AI developer services

Our research shows notable adoption of cloud AI developer services (CAIDS) across cloud platforms. Organizations configure these services to build tailored AI applications more efficiently by reducing the burden of managing infrastructure and leveraging ready-to-use tools and scalable resources.

Of the workloads we analyzed, 60% of the organizations using Microsoft Azure have configured Azure Cognitive Services (a collection of different AI services including Azure OpenAI Service), 40% have configured Azure Machine Learning workspaces and 28% have configured the Azure AI Bot Service. Among organizations using AWS, 25% have configured Amazon SageMaker and 20% have configured Amazon Bedrock. Among organizations using GCP, 20% have configured Vertex AI Workbench notebook instances.

These configuration rates indicate that organizations are deploying and fine-tuning AI models, with cloud platforms serving as the foundation.


## Why are Microsoft Azure's AI services more widely configured than the others? A few possible reasons:

- Organizations already using Microsoft products may find Azure AI services integrate well with existing infrastructure and operational needs. They may be able to streamline identity and access management through Microsoft Entra ID (formerly Azure Active Directory) and AI development through integrated tools like Azure DevOps.
- Another key driver could be Azure's hosted platform for OpenAI models, widely recognized as a leading solution for AI workloads.

## Configuration rates of AI services - % of organizations

### Microsoft Azure AI services

60% **Azure Cognitive Services**




60% of organizations that are using Microsoft Azure have Azure Cognitive Services configured

40% **Azure Machine Learning**



40% of organizations that are using Microsoft Azure have Azure Machine Learning configured


28% **Azure AI Bot Service**



28% of organizations that are using Microsoft Azure have Azure AI Bot Service configured


### AWS AI services

25% **Amazon SageMaker**



25% of organizations that are using AWS have Amazon SageMaker configured


20% **Amazon Bedrock**



20% of organizations that are using AWS have Amazon Bedrock configured

### GCP AI service

20% **Vertex AI Workbench in Google Cloud**



20% of organizations that are using GCP have Vertex AI Workbench notebook instances configured

Learn to set realistic cloud security goals using our cloud security maturity model

[Download now](#)





# AI risks and misconfigurations in the cloud

According to [OWASP](#), attackers seek to compromise AI models, manipulating their input data and the outputs they produce, exposing sensitive information and causing models to behave in undesirable ways. Overprivileged identities, exposed storage buckets, lack of auditing and lack of encryption are just some of the misconfigurations that open the door to such compromise.

Note that despite security measures cloud providers apply to the infrastructure and services

they provide, and their [playbooks for minimizing risk](#), managed cloud AI services can still introduce risk. For example, cloud provider defaults, designed for a seamless user experience, are often excessively permissive and left by users configured as such. Cloud customers are responsible for the security of the applications they deploy through managed cloud services, in accordance with the [shared responsibility model in the cloud](#).

We found evidence of the following risks in the cloud AI environments observed.



## Unremediated critical vulnerabilities

Over two thirds (70%) of the cloud workloads we analyzed that have an AI package installed have a critical vulnerability, compared with 50% of cloud workloads without AI installed.

For example we observed [CVE-2023-38545](#), a critical curl vulnerability, in more than one third of the cloud AI workloads analyzed. As of November 2024, the end of our data collection period, this vulnerability remained unremediated — more than a year after the CVE was published. When exploited the vulnerability can lead to unintended access to a rogue server.

One possible reason for the higher incidence of critical vulnerabilities is that many AI workloads run on Unix-based systems, which run many libraries, including open source, and for which vulnerabilities are often reported. When AI

assets are vulnerable, the outcome of exploitation is riskier due to potential manipulation of AI models, data tampering and leakage. If the vulnerable AI workload is also publicly exposed — creating a toxic combination — the likelihood of compromise increases significantly.

The high incidence of a critical vulnerability — at least one; there may be more! — in AI workloads is concerning due to the potential sensitivity of the data in the workload. Training and testing data may contain real information such as personal information (PI), personally identifiable information (PII) or customer data related to the nature of the AI project. Vulnerability management and security teams must be strategic about mitigating vulnerabilities, especially in AI workloads.


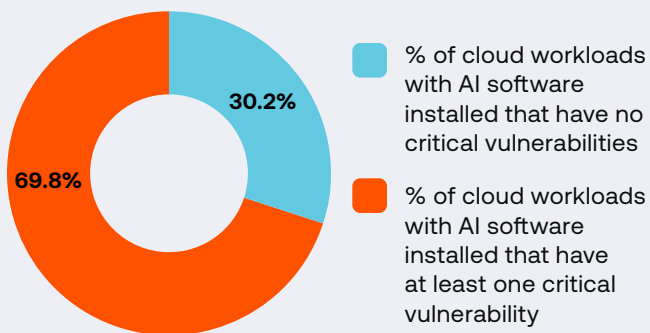


Figure 3 - At least one critical vulnerability was found in 69.8% of cloud workloads with AI software installed, compared with 50.1% of cloud workloads without AI software installed

## Jenga concept meets AI


More than three quarters (77%) of the organizations we studied have the overprivileged default Compute Engine service account configured in GCP Vertex AI Workbench notebook instances. We arrived at this finding after checking if the problematic Jenga concept also exists in AI services.

In Google Cloud, when a user creates a virtual machine, the Compute Engine service requires that a service account be attached. We knew that, by default, this attached service account has the Editor role – providing broad access to resources in the project. With this risky default in mind, we looked at GCP Vertex AI Workbench, a Jupyter notebook-based environment that provides a managed platform for developing machine learning workflows.

We found that each time a user creates a Vertex AI notebook instance, GCP creates – behind the scenes – a Compute Engine instance within the user's project. From previous observations we knew that when creating a virtual machine manually users are more likely to follow least privilege best practice in configuring the required attached service account.

In contrast, we found that when the instance is created as part of a notebook setup, 77.1% of organizations have the overprivileged default Compute Engine service account attached in at least one notebook. Jenga surfaces in AI: the underlying Compute Engine's overprivileged default configuration puts Vertex AI notebook instances at risk – a grave concern especially in AI systems handling sensitive data.

The Jenga concept, introduced by Tenable Cloud Research, describes the tendency of cloud providers to build one service on top of another, with any single misconfigured service putting all the services built on top of it at risk. Users are largely unaware of the existence of these behind-the-scenes building blocks as well as of any risk propagated from inherited defaults.



- % of organizations with Vertex AI Workbench installed that do not have default service account risk
- % of organizations with the overprivileged default Compute Engine service account attached in at least one Vertex AI Workbench notebook

Figure 4 - 77.1% of the organizations that have GCP Vertex AI Workbench configured have at least one notebook instance configured with the overprivileged default Compute Engine service account

# Amazon Bedrock training bucket without public access blocked

Among the organizations that have configured Amazon Bedrock training buckets, 14.3% have at least one bucket that does not have Amazon S3 Block Public Access enabled.

The Amazon S3 [Block Public Access](#) feature, considered a best practice for securely configuring sensitive S3 buckets, is designed to prevent unauthorized access and accidental data exposure. However, we identified instances in which Amazon Bedrock training buckets lacked this protection — a configuration that increases the risk of unintentional excessive exposure. Such oversights can leave sensitive data vulnerable to tampering and leakage — a risk that is even more concerning for AI training data, as [data poisoning](#) is highlighted as a top security issue in the [OWASP Top 10 threats for machine learning systems](#).

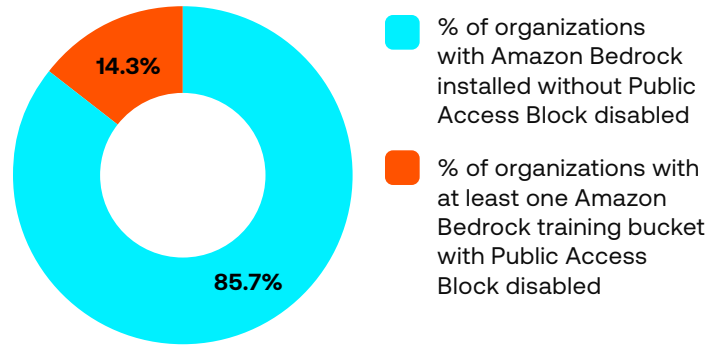


Figure 5 - Potentially risky public exposure was found in at least one Amazon Bedrock AI model training bucket — which can contain sensitive data — in 14.3% of organizations that have configured the Bedrock service

## AI big data has big implications for cloud workload security

Cloud AI services differ from traditional cloud services not only in scale but by requiring iterative data processing for training, inference and predictive analysis.

Due to the enormous amount of data involved, cloud AI workloads have a higher chance than standard workloads of containing sensitive data, heightening security risks when misconfigured. Configuration flaws can lead to corrupted models, inaccurate predictions and disruptions in AI-based processes — and exploitations such as unauthorized access, data exfiltration and AI model data tampering, as well as compliance violations.

Addressing these risks is vital to safeguarding AI applications in the cloud.



## Amazon Bedrock training buckets are overly permissive

A small but important portion (5%) of the organizations we studied that have configured Amazon Bedrock have at least one overly-permissive training bucket.

Overly-permissive storage buckets are a familiar cloud misconfiguration; in AI environments such risks are amplified if the buckets contain sensitive data used to train or fine-tune AI models. If improperly secured, the overly-permissive buckets can be compromised by attackers to modify data, steal confidential information or disrupt the training process.

We examined the policies of Amazon S3 buckets used to train AI models in Amazon Bedrock environments. We determined which buckets are overly permissive, according to the permissions granted in the policy and the permissions in use, and identified bucket policies that do not align with least privilege best practice.

The significance of these findings cannot be understated. For an attacker with access to the environment through a prior breach or public exposure, over-permissiveness is delicious candy at the supermarket checkout counter. Upon a breach, an attacker could — even if previously unaware of them — discover and exploit the excessive permissions to compromise training buckets and steal proprietary AI training data. The reputational and financial impact to the organization could include the loss of competitive advantage inherent in the future AI application.





Figure 6 - 4.8% of organizations that are using Amazon Bedrock have at least one overly-permissive AI training bucket

## Amazon SageMaker with root access enabled

The vast majority (90.5%) of organizations that have configured Amazon SageMaker have the risky default of root access enabled in at least one notebook instance.

By default, when a notebook instance is created, users who log into the notebook instance have root access. Granting root access to Amazon SageMaker notebook instances introduces unnecessary risk by providing users with administrator privileges. With root access, users can edit or delete system-critical files, including those that contribute to the AI model, install unauthorized software and modify essential environment components, increasing the risk if compromised. According to [AWS](#), “In adherence to the principle of least privilege, it is a recommended security best practice to restrict root access to instance resources to avoid unintentionally over provisioning permissions.”

Failure to properly adhere to the principle of least privilege significantly increases the risk of unauthorized access, enabling attackers to exfiltrate AI models — that is, steal models that may expose proprietary algorithms and intellectual property. As mentioned, compromised credentials can allow attackers access to other critical resources, such as S3 buckets, which often store training data, pre-trained models or sensitive information such as PII. The consequences of such breaches are severe.



- % of organizations with Amazon SageMaker configured without root access risk
- % of organizations with Amazon SageMaker configured that have risky root access enabled

Figure 7 - The risky default of root access enabled in at least one Amazon SageMaker notebook instance was found in 90.5% of organizations that have configured the service

# Mitigation strategies for AI risks

**Like Mickey Mouse learned as a sorcerer's apprentice: When deploying something with vast powers, rein in risk from the start.**

Like any component in a cloud environment, AI tools, models and data can introduce risk in the form of misconfigurations, risky entitlements and vulnerabilities. As businesses rush to implement AI, DevOps may find themselves pressured to move quickly — driven to download self-managed AI packages. Yet, these packages are commonly open-source, with no guarantee they are secure. Likewise, teams may configure managed AI services from their cloud providers without fully assessing them for vulnerabilities or excessive privileges (and failing to spot Jenga-style risky defaults).

## Here are some recommended AI risk mitigation strategies to consider:

- ➔ **Manage exposure of your AI systems and data.** Take a contextual approach to revealing exposure across your cloud infrastructure, identities, data, workloads and AI tools. Monitor all assets, and integrate telemetry and security configurations, on-prem and in the cloud. Unified visibility and prioritized actions across the attack surface enable managing risk as environments change and AI threats evolve.
- ➔ **Classify all AI components linked to high-business-impact assets as sensitive** (e.g. sensitive data, privileged identities). Include AI tools and data in your asset inventory, scanning them constantly and understanding the risk if exploited. Even test data can be sensitive; if leaked, the risk is as for production data.
- ➔ **Keep pace with emerging AI regulations and guidelines.** Stay compliant by mapping key cloud-based AI data stores and implementing required access controls. Ensure your AI engineers are carrying out “secure by design” development, deployment and use of your organization’s AI systems, and following [NIST guidelines](#).
- ➔ **Apply cloud provider recommendations for their AI services.** Follow any playbooks for avoiding risky configurations — while bearing in mind that defaults are commonly insecure and such guidance is still evolving. In deploying services, ensure that any resource provisioned in the process adheres to best practices and the principle of least privilege.
- ➔ **Prevent unauthorized or overprivileged access to cloud-based AI models/data stores.** Reduce excessive permissions and tightly manage cloud identities using robust tools for least privilege and security posture. These should also detect and Jenga-style misconfigurations in your cloud AI environment.
- ➔ **Prioritize vulnerability remediation by impact.** Understand which CVEs have the greatest risk severity in your environment. Less sophisticated cloud security solutions can bombard with team notifications; advanced tools improve remediation efficiency and effectiveness, and reduce alert fatigue.

While the scope of our research does not include shadow AI risk, we are aware that it is a challenge for many organizations. It’s important to minimize shadow AI risk via centralized governance, education and monitoring. Protect your organization even further with policies that disallow unsanctioned AI applications, and educate employees on the risks, responsible AI use and approved alternatives. Monitor your cloud environment for shadow AI, limiting access as needed.

Securing your cloud AI workloads requires a robust, AI-conscious cloud native application protection platform (CNAPP) that defines sensitivity and contextualizes risk across your cloud infrastructure, workloads, network, identities, data and AI resources, enabling teams to expose risk and prioritize remediation amid an expanding attack surface. For broad visibility and effective security seek a comprehensive solution that spans cloud and on premises environments. Taking a hybrid-cloud security approach pushes the bar on insight into how AI exposure, or AI risk scoring, bubbles up and into your organization’s risk rating index.

# Closing thoughts

As the research shows, organizations are actively introducing AI in their development environments as AI frameworks and services make the effort so much easier. Increased use of AI creates vastly higher volumes of data for an organization, making the cloud — excellent at handling dynamic data stores — a natural AI growth platform.

But cloud-based AI has its security pitfalls. AI services and frameworks are commonly misconfigured. And AI components in the cloud often contain sensitive data, including intellectual property, proprietary algorithms and the AI models themselves, making them an attractive target for misuse and exploitation, and causing greater risk if not effectively secured. Despite this, most companies have mitigated only a small portion of their AI-related risks (Source: [Artificial Intelligence Index Report 2024](#), Stanford University).

Security leaders have the mandate and power to enable AI for their organizations, at minimal risk. It's a perfect time, at this early stage, to implement exposure management solutions and security best practices that turn AI aspirations into secure business benefits.

## Methodology

Tenable Cloud Research created this report by analyzing the telemetry gathered from workloads across diverse public cloud and enterprise landscapes, scanned through Tenable products (Tenable Cloud Security, Tenable Nessus). The data was collected between December 2022 and November 2024.

The data set consisted of:

- Cloud asset and configuration information
- Real-world workloads in active production
- Data from AWS, Azure and GCP environments

## About Tenable Cloud Research

Tenable Cloud Research is the cloud security research arm of Tenable Research. It conducts ongoing research into new attack vectors, uncovers and discloses cloud provider vulnerabilities and applies its expertise to innovatively fortify the Tenable Cloud Security product with innovations against emerging risks. Recent discoveries and research publications include:

- [New Attack Techniques in OPA and Terraform](#)
- [CVE-2024-8260: SMB Force-Authentication Vulnerability in OPA](#)
- [CloudImposer: RCE Vulnerability in GCP Composer](#)
- [ConfusedFunction: Privilege Escalation Vulnerability](#)
- [Abusing Service Tags to Bypass Azure Firewall Rules](#)
- [FlowFixation: AWS Apache Airflow Service Takeover](#)
- [2024 Cloud Risk Report](#)

## About Tenable

Tenable is the exposure management company, exposing and closing the cybersecurity gaps that erode business value, reputation and trust. The company's AI-powered exposure management platform radically unifies security visibility, insight and action across the attack surface, equipping modern organizations to protect against attacks from IT infrastructure to cloud environments to critical infrastructure and everywhere in between. By protecting enterprises from security exposure, Tenable reduces business risk for approximately 44,000 customers around the globe. Learn more at [www.tenable.com](http://www.tenable.com).

## Contact Us

Please email us at [sales@tenable.com](mailto:sales@tenable.com) or visit [tenable.com/contact](http://tenable.com/contact).