

Tenable Cloud Security Risk Report 2025

Exposed sensitive data, poorly kept secrets, good news for identity and workload security – and mitigation insights



Table of contents

Executive summary	03
Key findings	04
Cloud data security — Sensitive data and secrets in danger	05
→ How we define data sensitivity	05
→ How much sensitive data is in the cloud — and how sensitive is it, really?	06
→ How we define secrets	06
→ Sensitive data: Stored in public locations, posing immediate risk	08
→ Secrets aren't safe	09
Cloud workload protection — Toxic cloud trilogies remain	10
Cloud identity security — IdPs help but do not guarantee security	11
AI security — High-stakes cloud security	12
Mitigation strategies	13
Conclusion	14
Methodology	14
About Tenable Cloud Research	14

Executive summary

The Tenable Cloud Security Risk Report 2025 explores a broad range of exposures in cloud environments that can potentially increase risk for organizations. We examine the concerning ways some organizations are storing sensitive data and secrets — such as encryption keys, tokens or access keys that authenticate, authorize and encrypt data. We explore the prevalence of toxic cloud trilogies — workloads that simultaneously have a critical vulnerability, excessive permissions and public exposure. We analyze trends we're seeing in identity provider (IdP) services and discuss how organizations might extend this best practice to further alleviate risks, including credential abuse. And we revisit key findings from our [Cloud AI Risk Report 2025](#) concerning the use of AI developer tools and services. Last, we offer mitigation advice to help organizations address these challenges in their cloud environments.

Tenable Cloud Research created this report by gathering and analyzing the telemetry from workloads across diverse public cloud and enterprise environments, scanned using the Tenable Cloud Security platform between October 2024 and March 2025. (The AI findings cited here are from data collected between December 2022 and November 2024.)

This report provides a deep dive into the most pressing cloud security issues observed, highlighting the areas of data, identity, workload and AI, and offers mitigation guidance for organizations seeking concrete actions to close and avoid such cloud exposures.

Key findings

Our analysis of cloud environments revealed key areas of risk as well as favorable indications of improved security hygiene and best practice implementation.

Join the [Tenable Cloud Security research team](#) for a webinar examining findings from this report.



9% of the publicly accessible storage analyzed contains sensitive data

Public and sensitive – a dangerous intersection. Across cloud providers, 9% of the publicly accessible storage resources analyzed contain sensitive data — with 97% of that data classified as restricted or confidential. This kind of exposure creates an ideal entry point for threat actors and poses a serious, immediate security risk. The convergence of public access and sensitive data underscores the hidden dangers that come with the ease, convenience and visibility challenges of modern cloud computing.



54% of the organizations using Amazon Web Services (AWS) ECS task definitions have at least one secret embedded in their configurations

Secrets aren't safe. Over half (54%) of organizations using Amazon Web Services (AWS) Elastic Container Service (ECS) task definitions have at least one secret residing there — creating a dangerous exposure path. A similar pattern appears among organizations using Google Cloud Platform (GCP) Cloud Run (52%), while Microsoft Azure Logic Apps workflows show a lower but still significant rate of 31%. Even more concerning, we found that a surprising 3.5% of AWS Elastic Compute Cloud (EC2) instances overall have a secret in their user data. Given EC2's widespread use, this small percentage represents outsized risk.



29% of organizations have at least one toxic cloud trilogy — an encouraging decline from the previously studied period

Good news for cloud workload protection. The number of organizations with a toxic cloud trilogy dropped by nine percentage points, from 38% in January-June 2024 to 29% in October 2024-March 2025. Yet, 29% is still concerning: toxic cloud trilogies — meaning a cloud workload that is publicly exposed, critically vulnerable and highly privileged — must be surfaced and addressed.



83% of organizations using AWS have configured IdP services

Good news for cloud identity security, too. 83% of the organizations using AWS have configured IdP services — a best practice for managing cloud identities, improving authentication and simplifying access control. But IdPs alone can't eliminate risk. Overly-permissive access defaults, risky entitlements and standing permissions still leave organizations vulnerable to serious threats.



77% of organizations setting up Vertex AI Workbench in Google Cloud misconfigure at least one notebook with an overprivileged default service account

An AI-risk finding worth repeating. The [Tenable Cloud AI Risk Report 2025](#) reveals that, as with traditional cloud services, AI services are constructed from Jenga®-like building blocks that inherit hidden risky default configurations. One striking example we found is that 77% of organizations that had set up Vertex AI Workbench in Google Cloud had at least one notebook instance configured with the overprivileged default Compute Engine service account. This misconfiguration enables an attacker who has already compromised the notebook to broaden the risk impact to the entire environment and other services.

See how [Tenable Cloud Security](#) can help you close critical cloud exposures fast, even if you only have five minutes.

Cloud data security — Sensitive data and secrets in danger

Sensitive cloud data needs to be protected.

To secure it, you need to see and know where sensitive data lives in your environment. With the rapid rise in AI adoption, this visibility is more critical than ever. Cloud-native AI services, model training workloads and related stored data all rely on large, often sensitive datasets — raising the stakes for cloud data security.

Our research found sensitive data — including secrets, which are among the highest-risk asset types — stored in locations where it shouldn't be, creating significant and avoidable risk.

How we define data sensitivity


Sensitive data is data that carries some level of risk. In our research, we examined data across four sensitivity levels defined by Tenable Cloud Security during its scans of environments and workloads. Simply put, any data classified as private, confidential or restricted poses significant risk if exposed.




Restricted. Sensitive regulated data. Data that is highly sensitive; typically involves personally identifiable information (PII), such as driver's license numbers and classified government information or other critical data. Unauthorized access or exposure can lead to significant legal, financial or reputational damage. This category also includes secrets (see sidebar, page 6).



Confidential. Sensitive intellectual property. Data that requires protection from unauthorized access, such as financial data and employee records. Access is limited to authorized personnel.



Private. Business sensitive (need to know). Internal data that is highly confidential and meant for use within the organization but not for public consumption, such as dates of birth and nationality. Unauthorized access can pose some risks.



Public. Can be shared publicly. Data intended for the public whose exposure has no adverse consequences.

Learn how to [automate discovery of sensitive data](#) across your multi-cloud environments

How much sensitive data is in the cloud — and how sensitive is it, really?

Finding: Across cloud providers, 9% of all analyzed storage resources contain sensitive data — almost all of which is classified as restricted or confidential.

We began by examining the most common data storage types in each cloud service provider (CSP), including AWS S3 buckets, Azure Blob Storage containers and Google Cloud Storage buckets. Our analysis revealed that approximately 9% of all storage resources contain sensitive data that could pose a security risk if accessed by unauthorized parties.




Figure 1: 8.6% of the storage resources contain sensitive data

How we define secrets

Secrets are privileged credentials used by human and machine identities to access sensitive systems. In today’s cloud environments these include API keys, access keys, encryption keys and tokens in addition to traditional usernames and passwords. The Verizon 2025 Data Breach Investigations Report (DBIR) highlights the use of these credential types for diverse needs, across cloud infrastructure, web applications, development pipelines (CI/CD) and databases. Verizon notes that given how routinely system administrators and developers use them, it’s “not surprising that these secrets sometimes accidentally end up in public code repositories” where, depending on configuration, they can grant attackers direct access to environments.



Next, we sought to understand the levels of sensitivity associated with the stored data we analyzed. We made some interesting discoveries: roughly one third (31%) of the storage resources containing sensitive data held information classified as restricted — the highest data sensitivity level in [Tenable Cloud Security](#) — while about two thirds (66%) held information classified as confidential, the second highest level. Only a small percentage of sensitive data storage resources (3%) fell in the private category, where the potential risk if exposed is lowest.

% of sensitive storage resources containing restricted or confidential sensitive data




Figure 2: 97.3% of the sensitive storage resources contain information classified as restricted or confidential — meaning most exposed sensitive data is highly critical, making even limited exposure potentially damaging.

When comparing the object storage services of cloud providers (AWS, Azure and GCP) we found that AWS has a higher percentage of identified sensitive storage resources by far, at 17%. This compares with only 3% of sensitive storage resources identified in Azure environments and 7% in GCP environments.




Figure 3: When comparing the percentage of identified sensitive storage resources across cloud provider object storage services, AWS has the highest percentage, at 16.7%.

One reason we’re seeing a comparatively higher percentage of sensitive storage resources in AWS may be that users are confident in the AWS security measures they have put in place. Another possibility is that, as the oldest major public cloud, AWS has simply accumulated more sensitive data over time.

Sensitive data: Stored in public locations, posing immediate risk

Finding: 9% of the publicly accessible storage resources across cloud providers contain sensitive data — a critical exposure that presents an immediate and serious security risk.

Stop right there, as this is a clear warning sign. Nearly one in 10 public storage resources analyzed held sensitive information. Such misconfigurations can have severe consequences, including customer data leakage, secrets theft and potential financial loss, including fines from lack of regulatory compliance.

% of sensitive storage resources in publicly accessible storage




Figure 4: Cause for concern: 9.2% of public storage contains sensitive data

What lies behind this finding? Are organizations unaware their data is publicly accessible or are they unaware that their data is sensitive? Two major causes of inadvertent exposure are misconfigured access settings and overly permissive policies. Another is privilege elevation by developers — intended for short-term use but frequently forgotten, eventually becoming permanent. Other contributing factors include flawed permission structures caused by inconsistent access policies or overlapping roles, inadequate (often manual) monitoring and even the false belief that obscure storage bucket URLs provide sufficient protection. A major reason for the presence of sensitive data in public storage may be that organizations are unaware of the sensitivity level of the data.

Interestingly, we also found that 9% of non-public storage resources contained sensitive data — the same percentage as in public storage. While it's the expected practice that sensitive data is stored in non-public locations, it's not necessarily safe there either — misconfigurations or weak access controls can still lead to exposure or compromise.

For clarity, a note on public vs non-public storage: Public storage refers to storage resources that can be accessed without authentication or without specific permissions requirements. Non-public storage refers to storage resources that can be accessed only with authentication and by identities entitled with the necessary permissions.

Secrets aren't safe

Finding: 54% of organizations with ECS task definitions have a secret stored in at least one of them, risking credential exposure that could lead to data breaches, significantly higher cloud costs or even full cloud environment takeover.

As bees to pollen, attackers are drawn to secrets. Compromising secrets is a common attack vector of bad actors, allowing them to move laterally, escalate their privileges and gain access to valuable assets. Verizon's [2025 Data Breach Investigations Report \(DBIR\)](#) confirms the scale of this risk, finding that cloud infrastructure secrets account for 15% of all exposed secrets in public code repositories — making it the third largest category of exposed secrets, after web app infrastructure (39%) and development and continuous integration/deployment (CI/CD) (32%).


Indeed, we found organizations' secrets scattered across various cloud resources and potentially exposed through misconfigurations and oversight. Most prevalent is the finding of a secret in at least one AWS ECS task definition (54% of organizations) and in GCP CloudRun environment variables (52% of organizations). Also notable is the finding of secrets in the user data of at least one AWS EC2 instance (26% of organizations). As the most foundational AWS service, EC2 represents a high-risk exposure surface. We also looked at Infrastructure as Code (IaC), where we found at least one secret among 9% of organizations; propagation of this insecure practice bears long term risks.

% of organizations with at least one secret in at least one cloud resource

54.1% of organizations - AWS ECS task definitions




26.2% of organizations - AWS EC2 user data



51.6% of organizations - GCP CloudRun environment variables



30.5% of organizations - Azure Logic Apps workflow configurations



8.9% of organizations - Infrastructure as Code (IaC)




Figure 5: Among organizations, the most prevalent secrets are in AWS ECS task definitions (54.1%) and GCP CloudRun environment variables (51.6%).

Attackers are known to [penetrate cloud environments to leverage their compute power](#) for activities such as crypto mining — it is this scale of abuse that can cost companies huge amounts of money without their knowledge.

We looked at cloud environments across organizations to see in which cloud resources secrets exist. We also looked across cloud provider environments to identify cloud resources in which secrets are found.


The finding that 3.5% of AWS EC2 instances overall contain a secret in their user data is particularly concerning. EC2 user data is often used in automated configuration tasks and may contain a secret for purposes of authentication. However, storing secrets in plain text is poor practice, as a threat actor with initial access can use them to trigger a cascade of exploitative activity. Given AWS EC2's widespread use, such secrets are likely to exist in many instances in an environment.

Secrets can be leaked through many means, including git repositories, public storage and logs. Understanding where secrets reside, how they are used and by whom, is crucial for effective cloud risk management. Under [shared responsibility](#) guidelines, it falls to the organization — not the cloud provider — to ensure that secrets are properly secured.


➔ [Get pro-tips](#) for detecting, prioritizing and remediating common AWS misconfigurations and threats

% of cloud resources with at least one secret - across cloud providers


3.5% of AWS EC2 instances, in the user data



10.4% of AWS ECS task definitions



4.4% of Azure Logic App workflow configurations



9.3% of GCP Cloud Run environment variables




Figure 6: The most concerning finding across cloud providers is the percentage of secrets in AWS EC2 user data (3.5%), as EC2's centrality to almost every AWS architecture makes exposure highly likely.

Cloud workload protection — Toxic cloud trilogies remain

Finding: 29% of organizations have a toxic cloud trilogy, a decline of nine percentage points from the 38% found in our 2024 Cloud Risk Report.

Now for some good news: Significantly fewer organizations in this reporting period (29%) had at least one toxic cloud trilogy, compared to the 38% who had one in our 2024 Cloud Risk Report.

A toxic cloud trilogy is a cloud workload that is publicly exposed, critically vulnerable and highly privileged, creating a high-risk attack path and a prime target for bad actors.

Why the reduction in the prevalence of toxic cloud trilogies? Organizations may be increasingly leveraging cloud security solutions to prioritize risks by [impact and likelihood](#), focusing on the exposures that create high-risk attack paths. This includes mitigating critical vulnerabilities, public exposure and overprivileged human and service identities. Greater organizational efforts in cloud security education, training and guidance may also be helping. Nonetheless, these findings show that toxic cloud trilogies continue to pose an urgent problem for organizations.




Figure 7: Good news — we revealed a drop of nine points in the percentage of organizations (AWS/GCP) with at least one toxic cloud trilogy — but 28.6% is still cause for concern.



Cloud identity security — IdPs help but do not guarantee security

Finding: The vast majority of AWS users have configured IdPs, a best practice for managing identities; specifically, 83% of AWS organizations have configured at least one IdP.

Compromised identities remain a leading cause of cloud breaches so securing cloud identities is a critical priority. Using IdPs makes good sense as they help organizations centralize identity management, strengthen access control and authentication, and contend with multi cloud environment complexities.

Our research found that the vast majority of organizations using AWS have configured IdPs. This finding is reinforced by a recent Cloud Security Alliance (CSA) survey showing 75% of organizations are using two or more identity providers (CSA, October 2024, [“State of Multi-Cloud Identity Management Survey Report”](#)).

While IdPs support both human and service identities, their greatest value lies in managing human identities. Features like single sign on (SSO), multi-factor authentication (MFA) and conditional access policies help protect against credential abuse, phishing and other threats that target humans. Credential compromise not only enables initial access — it often leads to privilege escalation, lateral movement and access to sensitive data. As AI workloads grow, securing identities will be even more critical to maintaining access control and warding off increasingly sophisticated attacks. The Verizon 2025 DBIR reports that credential abuse remains the most common known initial access vector, implicated in 22% of breaches.

Given the growing popularity of IdPs, it's important to know that simply using them does not guarantee security. Overly permissive defaults, excessive entitlements and standing permissions can leave organizations exposed to identity-based threats. While MFA is proven to block unauthorized access, a recent [451 Alliance study](#) found its adoption is often hindered by user experience challenges. To fully realize the security benefits of IdPs, cloud security stakeholders must go further: enforce MFA, implement least privilege and actively manage identities and entitlements to reduce risk across their cloud environments.

% of AWS organizations that have configured IdPs




Figure 8: 83.4% of organizations using AWS have configured at least one IdP — a cloud security best practice that improves access control and authentication.

[Get best practices](#) for securing cloud identities, removing excessive permissions and enforcing just-in-time access.



AI security — High-stakes cloud security

Finding: 77% of the organizations that have GCP Vertex AI Workbench set up have at least one notebook configured with the overprivileged default Compute Engine service account

Finding: 70% of cloud AI workloads contain at least one unremediated critical vulnerability

From a cloud perspective, the exciting news about AI is that the cloud is AI's natural home. The cloud makes it much easier and manageable to handle the enormous compute and data load required to test and train AI models, and run generative AI workloads. On the other hand if compromised, AI workloads could have a profound impact not only on an organization's cloud environment but its business as a whole.

In our [Cloud AI Risk Report 2025](#) we surfaced two concerning AI risk factors in developer tools and services that bear repeating:

- ➔ **Overprivileged default service accounts.** 77% of the organizations that had GCP's Vertex AI Workbench set up had at least one notebook instance configured with the overprivileged default Compute Engine service account. Misconfigured defaults in AI service building blocks create high-risk exposure paths that can lead to privilege escalation, lateral movement and compliance violations. Such risks are amplified as AI grows as an attractive target for threat actors. In this case, the default service account had permissions beyond those needed for the AI service and notebook, increasing the risk impact across the broader environment and other services.
- ➔ **Workloads with critical vulnerabilities.** Our research found that 70% of AI cloud workloads across Azure, AWS and GCP had at least one unremediated critical vulnerability — compared with “only” 50% in non-AI workloads. A critical vulnerability in an AI environment can serve as a launchpad for unauthorized access to sensitive training data, model manipulation, data poisoning or even lateral movement within the broader cloud infrastructure. When part of a toxic risk combination, such as overly permissive access and/or public exposure, critical CVEs can amplify attacker success and persistence.

Given that the goal of many cloud breaches is to obtain sensitive data, and the massive quantity of AI data means a significant portion of it is likely to contain sensitive data, AI workloads need special security consideration.

Organizations using AI developer tools and services would do well to understand and mitigate cloud-based AI risks as early as possible in their development lifecycle. Security must be implemented in lockstep with an organization's AI initiatives. The good news is that the best practices for security cloud environments apply to securing AI environments.

Mitigation strategies

It takes a lot to shore up a cloud environment against a determined and highly motivated attacker. We suggest the following mitigating actions for the security threats identified in this report.

- ➔ **Monitor and minimize public exposure.** Not everyone managing cloud assets is familiar with secure storage practices, increasing the risk of unintended exposure. Continuously monitor for public access — including by third parties, often a weak link in the cloud security chain — and reduce sensitive data exposure by automating detection of misconfigured storage services, enforcing least-privilege and assessing posture on an ongoing basis. Use exposure management tools to map complex asset, identity and risk relationships across hybrid environments, to spot and prioritize cross-cloud attack paths.
- ➔ **Safeguard secrets through continuous visibility into where sensitive data resides.** Make secrets management one of the core pillars of your data governance strategy. The major CSPs offer mature, native secrets management tools that integrate easily with their identity and access management (IAM) frameworks: use them! Leveraging these tools is not just a best practice, it's essential to enforcing least privilege, reducing sprawl and improving auditability.
- ➔ **Prioritize vulnerabilities for remediation by combining context with likelihood of exploitation.** Correlate identity, vulnerability and network configuration data across your entire cloud stack to uncover toxic cloud trilogies — risky combinations that expose sensitive data and cloud infrastructure. Use vulnerability intelligence to assess the would-be risk impact and understand how specific exposures could affect your environment and business.
- ➔ **Secure your identities to secure your cloud.** Educate your IAM and security teams on the critical role of entitlements management in reducing excessive permissions. Build on your adoption of IdP — which CSPs have made easier to use — to take identity security one step further by implementing Just in Time (JIT) access to eliminate standing permissions and enforce timebound access. Seek out solutions offering [JIT for IdP groups](#) and that deliver via your go-to collaboration tools.
- ➔ **Secure your sensitive data in this age of AI.** Inventory, classify and track where your sensitive data resides across the cloud, including any AI or developer services that handle it. Know the sensitivity level and who has access when — so you have the context needed to apply the necessary controls to protect the data, and to understand and prioritize related risk.



Conclusion

With today's cloud environments offering fertile ground for attackers, automating risk management across cloud infrastructure, workloads, identities, storage, data and AI resources is essential. A mature cloud security platform (CNAPP) integrates with cloud-native tools, IdPs and collaboration platforms to reveal and prioritize risk — from secrets exposure and data sensitivity to CVEs and access misconfigurations — empowering security teams to stay focused, effective and ahead of evolving threats.

See how Tenable Cloud Security CNAPP can help you close critical cloud exposures fast — even if you have only five minutes — with a unified, comprehensive solution that automates discovery across identities, workloads, configurations, data and AI resources, enforces least privilege and prioritizes risks with context for speedy remediation — and integrates with Tenable One to offer hybrid cloud security.

Methodology

Tenable Cloud Research created this report by analyzing telemetry gathered from workloads across diverse public cloud and enterprise landscapes, scanned through Tenable Cloud Security. The data was collected from October 2024 through March 2025.

The data set consisted of:

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

Findings greater than X.5 have been rounded up to the next whole number; those less than X.5 have been rounded down.

The AI findings cited in this report are from a previous report whose data was collected between December 2022 and November 2024.

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.

Contact Us

Please email us at sales@tenable.com or visit tenable.com/contact.

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 fortify the Tenable Cloud Security product with innovations against emerging risks. Recent research publications and discoveries include:

- [Cloud AI Risk Report 2025](#)
- [Cloud Risk Report 2024](#)
- [ConfusedComposer: A Privilege Escalation Vulnerability Impacting GCP Composer](#)
- [ImageRunner: A Privilege Escalation Vulnerability Impacting GCP Cloud Run](#)
- [The Dark Side of Domain-Specific Languages: Uncovering New Attack Techniques in OPA and Terraform](#)
- [CVE-2024-8260: SMB ForceAuthentication Vulnerability in OPA Could Lead to Credential Leakage](#)
- [CloudImposer: Executing Code on Millions of Google Servers with a Single Malicious Package](#)
- [ConfusedFunction: A Privilege Escalation Vulnerability Impacting GCP Cloud Functions](#)
- [These Services Shall Not Pass: Abusing Service Tags to Bypass Azure Firewall Rules \(Customer Action Required\)](#)
- [FlowFixation: AWS Apache Airflow Service Takeover Vulnerability and Why Neglecting Guardrails Puts Major CSPs at Risk](#)

JENGA® IS A REGISTERED TRADEMARK OWNED BY POKONOBÉ ASSOCIATES.

COPYRIGHT 2025 TENABLE, INC. ALL RIGHTS RESERVED. TENABLE, NESSUS, LUMIN, ASSURE, AND THE TENABLE LOGO ARE REGISTERED TRADEMARKS OF TENABLE, INC. OR ITS AFFILIATES. ALL OTHER PRODUCTS OR SERVICES ARE TRADEMARKS OF THEIR RESPECTIVE OWNERS.

061125

