Nightfall AI

2025 State of Secrets Exposure Report

Table of Contents

Executive Summary

Purpose for the 2025 SOS Report

Nightfall's annual State of Secrets Report is designed to help security, IT, and compliance professionals make informed decisions about their data loss prevention (DLP) and data security strategies. Given the speed at which new technologies are transforming the ways businesses can leverage data, this assessment needs to be ongoing.

2024: The Year of Nonhuman Identities

Interestingly, many security teams were so focused on the critical security and privacy concerns of leveraging custom AI to drive new efficiencies last year that they missed the risks being created during the routine processes surrounding these efforts. For example, when rolling out aggregated data lakes, analysis tools, and AI models, engineers need to access and share nonhuman identities–tokens, keys, etc.–to each of these environment.

While much effort was put into ensuring sensitive data didn't enter these models unsecured, the keys that could grant attackers access to the data lakes, integrations, and models were less protected in many cases. So, it wasn't necessarily major AI models we saw putting organizations at grave risk last year–it was everything else.

Key Takeaways

44%

This is the percentage improvement we saw in risky password sharing practices. Read on to **find out why.**

Cloud Infrastructure API Keys

These are the most frequently shared secrets. You'll learn **what else topped the list** and what it means.

96.5%

of the secrets found were **API keys and passwords**. (We'll also tell you the ironic key types discovered.)

Key Background Information

What are *Secrets*?
Secrets are nonhuman and human identifiers including:

- **API keys**
- **Cryptographic keys**
- **Tokens**
- **Database connection strings**
- **Passwords**

Why are *Nonhuman Identities* referred to as Secrets?

Non-human identities are commonly called secrets because, like top secret information, they are valuable for their function. Secrets are authentication credentials like API keys, passwords, and certificates. These secrets enable non-human entities such as servers and applications to identify themselves and interact with other systems, in contrast to human identities that typically leverage user-managed username/password combinations.

Why are Secrets considered sensitive information?

Since these secret credentials are essential for system operations and their compromise can create serious security vulnerabilities, protecting them from unauthorized access is crucial. Further, it can be difficult to monitor or trace malicious activity back to a specific user or actor, making nonhuman identities extremely useful for attackers. This makes secret management and security a critical priority for cybersecurity teams who must ensure these non-human authentication mechanisms remain confidential and secure.

What are *Nonhuman Identities*?

A non-human identity (NHI) is a digital ID given to automated systems like machines, apps, or devices. NHIs allow these systems to access networks and data without human involvement–similar to usernames and passwords that grant access to private accounts for end-users, but designed for automated processes like APIs and IoT devices.

20X*

Over 360 security and IT professionals surveyed by ESG said that for every human access identity they manage, they have 20 times more NHIs to manage.

52%*

Of those surveyed expect managed NHIs to increase by 20% in 2025. This issue is not going away.

*Source - 2024 Enterprise Strategy Group (ESG) NHI Research e-Book

Detection Methodology

Secrets detection model design determines accuracy and results.

Detection Overview

Engine Capabilities

Nightfall's cloud native, AI-powered detection engine can ID data types that stump even the most popular solutions on the market. Models must be built, trained, and maintained by AI experts for specific data types with challenging characteristics common to secrets.

Accuracy Matters

When it comes to DLP, the first step is discovery. You need to find all the locations where sensitive data lives–every file, folder, app, post, message, and app. Since most modern organizations today leverage cloud storage, workspaces, and cloud-native SaaS apps, the possibilities of where data could exist have definitely widened. It could exist in unexpected storage locations, users' emails, direct messages, help desk applications, project management apps, endpoints, shared files, and even external sharing websites. With so many possibilities for places data can end up, discovery has proven the most difficult step in DLP.

Nightfall's Detection Model

2X More accurate

4X Fewer False Positives

If things break down at the detection / discovery phase, all the steps that follow are nullified. For example, how helpful is it to monitor security settings for sensitive files if you don't know what content is inside the file, or all the locations it might be stored? And what use is it to alert on unsanctioned sharing of sensitive data if you are unable to accurately differentiate between sensitive and nonsensitive data? Security teams may develop a sense of overconfidence, feeling their data remains secure, simply due to a lack of alerts. In reality, the data could be quite vulnerable, but a detection tool was unable to catch it. That's problematic at best.

Challenges in Building Data Detection Models

Early Points of Failure in Many Models

When classifying data as sensitive or nonsensitive, the first step is normalization, the process of converting each secret into plain text so it can be analyzed. Accuracy at this point is essential to support the entire workflow. If a model alters, misreads, or misinterprets data during this step, its decision-making after that point will be based on incorrect information.

Some Data in Structured Files Can be Challenging to Normalize

When detecting secrets in a CSV / tabular file / Google Sheet, many tools inspect only comma-separated values to define whether the file contains secrets or not. However, this can be misleading, as sensitive data is often high entropy in format. A *high entropy data format* refers to a data structure where the information is highly randomized or unpredictable, essentially means there is a large degree of variation and uncertainty within the data, often indicating a high level of complexity. This is commonly seen in encrypted data, for example, where the pattern is difficult to discern without the decryption key.

A high entropy data string *could* be a password, but it could also be a column of nonsensitive transaction IDs. The result of alerting on non-sensitive data is a noisy DLP that wears security teams out with false positives, causing analysts to mistrust and possibly even ignore alerts in the future. Teams simply don't have time to constantly sift through false positives to find the real events.

The Challenge of Unstructured Data

80-90% of corporate data is unstructured (Gartner).

If sensitive data were structured and never stored outside of a designated spreadsheet cell, it would be much easier to pinpoint. But that's not the nature of how data ends up where it does. Truth is, people use data in all sorts of ways and locations to accomplish their daily work, and most often, sensitive data isn't formatted in perfectly clean ways or places that make it easy to detect.

Secrets are especially prone to unstructured–and unpredictable–formats. Nightfall utilizes layered, AI and ML driven, context-aware, and content-aware detection models. Additionally, our platform goes far beyond the existing label of a file to determine its level of sensitivity, performing deep inspection of the contents. Once labeled based on sensitivity, Nightfall is also able to provide security teams with file history, security settings, and a list of users or teams who have access to those files.

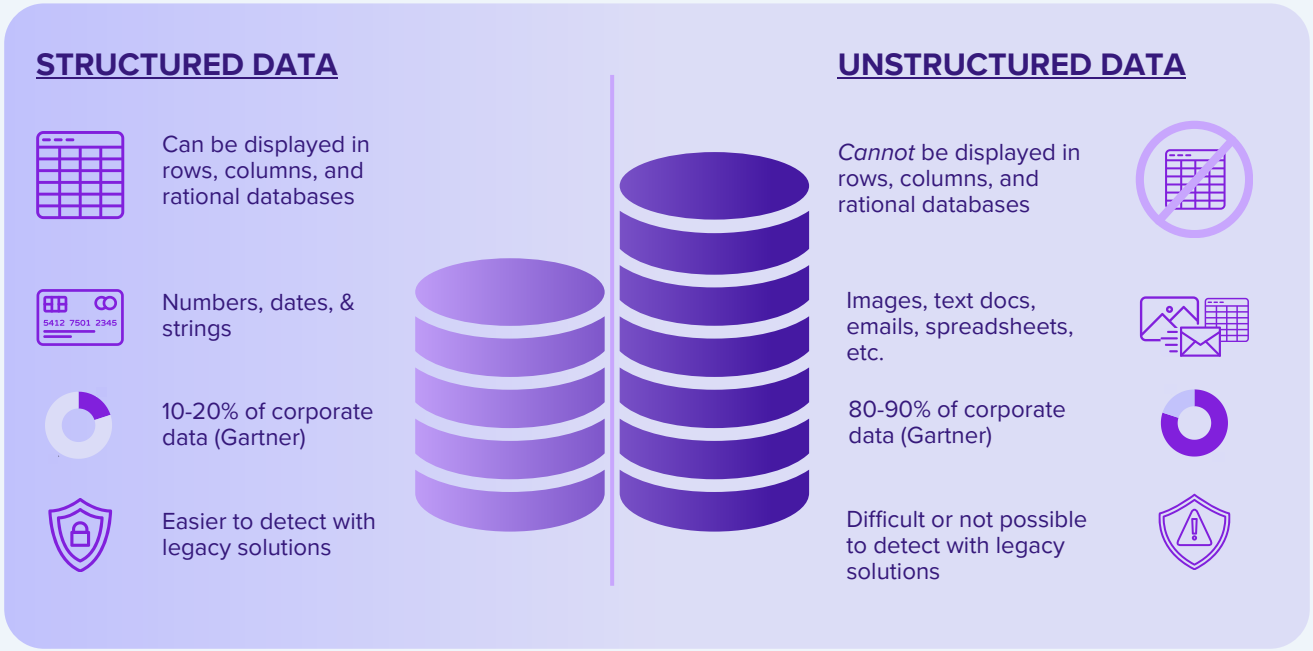

2025 Findings

What was shared, where, and how often last year.

2025 Findings: What We Found

Types of Secrets We Found Lurking In Systems Last Year

242.5K+

Total secrets discovered across SaaS apps and GenAI tools

API Keys took the most-shared lead in 2024. Of the 242,531 secrets discovered in SaaS last year, **API keys were the most-shared, surpassing exposed passwords for the first time.** Understanding sharing habits can help you know how to prioritize data security risks.

What Stayed the Same

Year over year, Passwords and API Keys are the two most shared secrets types. CISOs have been scratching their heads for years over password sharing–to little avail.

Security awareness training programs are not only cited as best practices (and have been for years), but they are required by many compliance frameworks.

A Microsoft study of security awareness training videos found they only improve rates of user errors by 3% at most, and that was specific to anti-phishing. Clearly a more effective means of training is needed to supplement employees' data security skills development.

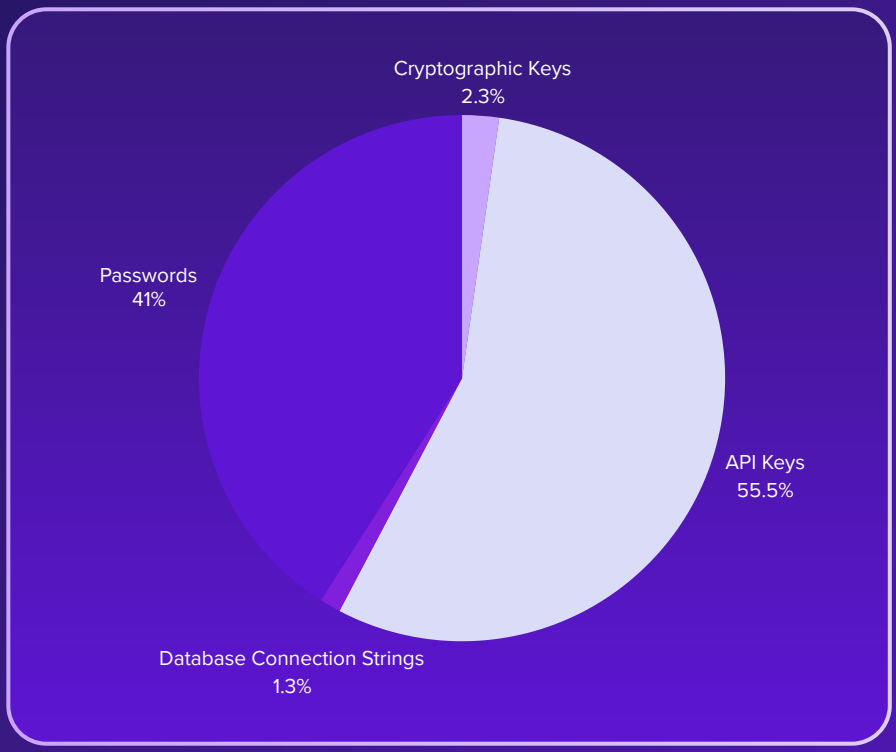

Key Takeaway

While Passwords are an issue everywhere, API Keys are likely to be among the top-shared secrets at organizations building software. The reason for this is that they are vital to productivity workflows in SaaS development. Organizations need a plan in place to find and remediate exposure of these two key data types.

2025 Findings: What Changed

1.3 API Keys and 1 Password Shared Per Week, Per Employee Last Year

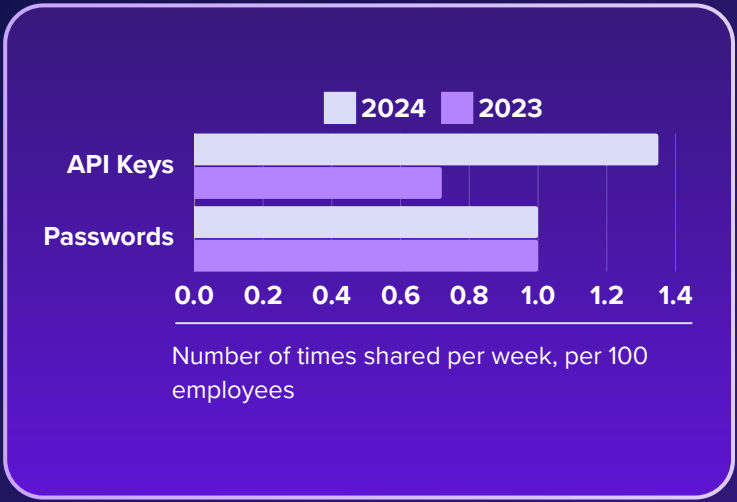

Over the last year, Nightfall users saw an improvement in employees' data hygiene habits. Their teams exposed **7% fewer API keys** and **44% fewer passwords per capita**. For comparison, the 2023 ratio of API keys to passwords was .72 keys for every 1 credential shared in unsanctioned SaaS locations. This year the ratio is reversed, with 1.35 API keys shared for every 1 password.

Implications of Improved Data Hygiene Practices

We're calling this a win for data hygiene habits among end users. Additionally, corporations are increasingly adopting password managers as part of their security stacks. Still, end users must choose to leverage them, bringing it back to improved practices. Good job, team.

One other factor is at play for Nightfall clients: a ***collaborative user-led remediation*** feature. This feature enables end users to fix their own errors within a time frame set by their security teams, allowing them to learn and take positive course correction on their own. Over time, habits improve. The significant improvement with password sharing points to some combination of improved habits and potentially better use of management tools.

Key Takeaway

The upside of the reversal is that what's changed is *not an increase* in API sharing rates. Rather, Nightfall users substantially improved their password sharing practices in 2024. That improvement over time, alongside collaborative remediation, password managers, and auto-remediation for events users miss seems to be **demonstrably improving habits while preventing breaches**.

2025 Findings: Top 3 API Key Types Found

Where's your biggest risk? 2025's Most-shared API Key Types

Understanding which key types teams are sharing most often helps you pinpoint where your biggest risk is, and which parts of your environment may need more attention.

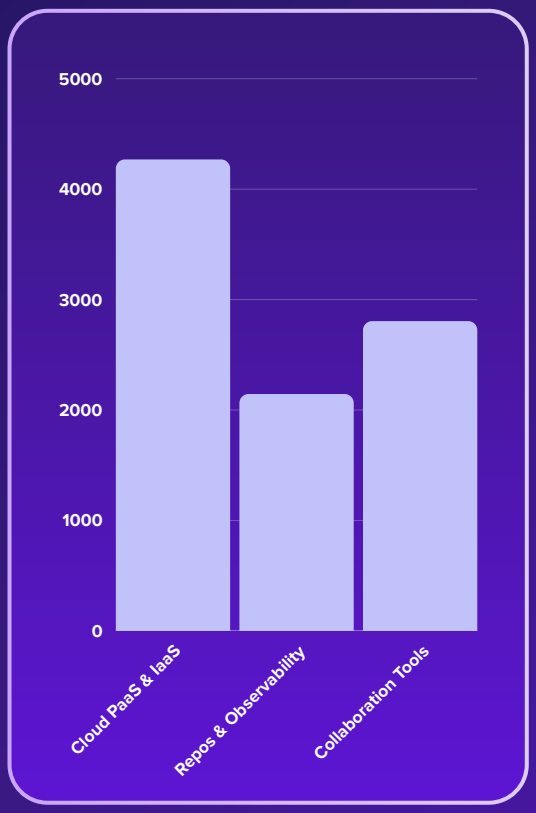

Cloud Infrastructure & Platform Keys

The **most frequently shared**–and arguably the riskiest–API keys were for cloud Platform as a Service (PaaS) and Infrastructure as a Service (IaaS) locations, like **AWS, Google Cloud Platform (GCP), and Azure**.

Implications of Cloud Infrastructure API Key Loss are significant. Attackers could gain unauthorized access to your most valuable digital resources, leading to breaches, manipulation of system configurations, use your servers to launch malicious attacks, etc. Learn more about the risks stolen API keys pose to your environment.

Developer Tool API Keys

While Developer Tool keys are not the second most shared keys, they pose the second greatest risk to your organization if uncovered by threat actors. In particular, **code repositories** ("repos") and **observability / monitoring platforms** can present serious issues if accessed by unauthorized users. If your teams are using these platforms to during the code review process, chances are the API keys are being shared in your systems during routine developer workflows.

Collaboration tool API Keys

API tools that grant access to collaboration tools like Slack, Confluence, and Notion pose far less risk, as their primary function is to create automated workflows or integrations. Certainly, they are not without unique risk, but despite the frequency with which these keys are shared there's far less that attackers can do with them. They're more likely to want to access the instances and accounts, themselves, to find other secrets and sensitive data. To do that, attackers will need to steal user credentials.

2025 Findings: Code Repository Keys

Implications of Code Repository API Key Loss

The frequencey with which code repo keys are shared to risky locations should give CISOs and engineering executives pause.

Common Code Repository Platforms

- GitHub
- GitLab
- Bitbucket
- Azure DevOps
- AWS CodeCommit
- SourceForge
- Apache Allura
- Gerrit
- Gitea
- Gogs
- Launchpad
- Phabricator
- Beanstalk
- RhodeCode
- Google Cloud Source Repositories

What if a code repository API key is stolen?
Most critically, attackers with repo API keys could:

Access and exfiltrate proprietary source code, potentially exposing:

- Trade secrets and intellectual property
- Security vulnerabilities that could be exploited
- Internal architecture details and business logic

Make unauthorized changes to the codebase:

- Inject malicious code or backdoors
- Sabotage critical systems
- Corrupt or delete repositories

Gain broader system access by:

- Finding hardcoded credentials in source code
- Discovering internal API endpoints and systems
- Identifying configuration details that could enable lateral movement

Disrupt development operations by:

- Blocking legitimate developer access
- Breaking build pipelines
- Tampering with deployment processes

Key Takeaway

Any cloud location used in your developer workflows can be compromised if attackers find API keys that grant them access to those environments. Unfortunately, **the top two riskiest API key types are also among the top 3 most-shared.** That creates a significant security challenge that needs prioritized attention.

2025 Findings: Rare But Financially Risky

Financial Platform API Keys

Each organization's most valuable digital assets will vary depending on its business model. For companies following a product-led growth (PLG) strategy in tech, **automated payment channels** are extremely valuable for revenue streams, making it easy to enter payment data. A compromised payment channel key could halt or redirect payments to unauthorized parties.

1.48%

of the API keys exposed in SaaS last year were for payment platforms like Stripe, Plaid, PayPal, and more.

Online Payment Platform Integrations in PLG-Driven Companies

Most startups prefer to minimize PCI DSS compliance exposure by using a payments vendor who offers an integrated experience for their buyers. These providers allow users to enter credit card information and move directly into the product they've just purchased–no sales process, just an immediate free trial. This has been a significant boon to rapid growth in tech, but creates one more API key type that's now floating around in organizations' environments.

Key Takeaway

Financial platform API keys are an immediate threat to revenue streams. An often overlooked threat vector, they're also not likely to receive specifically targeted protection unless your organization is actively mitigating data sprawl in daily employee workflows. Since human error is still a significant factor in data security events, we highly recommend implementing a powerful data detection and response tool as part of your cloud DLP strategy.

Risk to Your Organization

Fortunately, sharing rates for these platforms are somewhat lower. 1.48% of the API keys discovered last year were Financial. With these keys, an attacker could siphon off users' payments and direct them to their own accounts. In companies whose main method of new payment collection is online, this could lead to significant losses. What's more, detecting the problem likely wouldn't happen until someone notices a drop in payments flowing into bank accounts.

Even then, an investigation could take some time to understand the cause, as API keys are not associated with a user's credentials or other authentication measures that give context to changes or events. It's best to proactively remove them from systems.

2025 Findings: API Keys to Corporate LLMs

7.67%

of the API keys we found last year were for AI platforms.

The LLM Attack Vector People May Forget About

In the rush to protect data going into LLMs from compliance violations, many organizations forget about their GenAI API keys to tools like OpenAI, Anthropic, Hugging Face, or Cohere.

What's the risk to your organization of a stolen AI platform API key?

LLM use is helpful, and not just for legitimate corporations. It's also really useful to attackers who may have any number of malicious uses for these tools. Thing is, attackers are not known for their generous willingness to pay for quality services. If they want it, they're going to go steal it. With your API key, they can leverage your corporate LLM instance without paying a dime. Lose these, and you could inadvertently be funding criminal activities. And since AI use is expensive, you're going to lose a bit of cash in the process.

Key Takeaway

LLMs are still new for most organizations, with adoption rates skyrocketing. Attackers are likely to bet on security gaps (such as a loose API key that gives them access to your corporate instance). Alongside efforts to ensure compliant use of data within corporate LLMs, it's important to secure their API keys to avoid cost overages due to resource theft.

2025 Findings: Identity + Access Management

It only helps if it's secured.

Of all the API key types we found in systems last year, this category is the most ironic. Access management tools are purchased in order to protect one's environment from cyber attack. The two specific categories discovered include **Identity Access Management** and **Security Operations** tools.

6.36%

of API keys discovered this year were specific to identity and identity access management tools.

Key Takeaway

Access management keys interact with a cloud platform's identity system. It lets you programmatically control user accounts, set permissions, and define who can access different resources and perform specific actions through the framework.

Type and frequency of API keys found

Identity Management: 6.36%
(Okta, Google Cloud IAM, Microsoft Entra, authO, jwt, etc.)

What Attackers Can do With Identity Tool APIs:

Critical, Immediate Risks

- **Identity theft and account takeover** - Attackers can impersonate legitimate users and administrators, potentially gaining broad system access.
- **Privilege escalation** - If the stolen keys have administrative capabilities, attackers can modify permissions and create new privileged accounts
- **Data exfiltration** - Unauthorized access to sensitive systems and data protected by IAM controls
- **Backdoor creation** - Attackers can establish persistent access by creating new credentials or modifying authentication rules
- **Lateral movement** - Using compromised identities to pivot across systems and escalate privileges
- **Audit trail manipulation** - Activity using stolen valid credentials may appear legitimate, making detection harder
- **Authentication bypass** - IAM keys could be used to disable or weaken existing security controls
- **Business Impacts -** A compromise of these keys can result in **regulatory fines** due to compliance violations, **service disruptions** if active keys need to be revoked, and **financial losses** due to incident response and recovery.

2025 Findings: *Everything Else*

Other API Key Types in Systems

Numerous other types of API keys were found in systems last year. While they are comprised of less frequently shared key types, each holds unique risk if found by bad actors. For example, some organizations use file sharing locations to send and receive sensitive information from their MSSP providers, like penetration tests and security audits. That risk is going to be very different from the risk to customer data if a CRM API key is discovered by an attacker. That kind of access could cause a very expensive breach, including compliance fines, legal fees, and reputational damage.

Types and frequencies of API keys found

Security Operations Tools: 0.02%
(SIEM, MDR, DDR, SSMP, etc.)

File Storage and Sharing: 0.02%
(Box, Dropbox, etc.)

Customer Relationship Management (CRM): 0.03%
(Salesforce, HubSpot, Zoho, Monday.com,etc.)

Cybersecurity Tools: 0.02%
(SIEM, EDR, DLP, etc.)

Social Media Accounts: 1.75%
(Facebook, X, Telegram, etc.)

Email Services: 0.32%
(Mail Chimp, Sendgrid, etc.)

Key Takeaway

t's also important to remember that while some API key types are shared at a much lower frequency, **it only takes one**.

In Summary

Secrets that remain exposed in vulnerable SaaS locations create a tremendous amount of risk. In a world where attackers are constantly at the door and will likely make it in at some point, the best thing you can do is ensure there's nothing for them to find. Continuously scanning to identify, classify, and remove sensitive data from systems is an essential component of an effective data loss prevention (DLP) strategy. Secrets are meant to be kept.

Risky Business

What leaked secrets could mean for your organization.

Security Risks of Compromised Secrets

Exposed Cloud PaaS or IaaS API Keys

Depending on just what's *in* your PaaS or IaaS, the damage that can occur due to an exposed API key granting access to those locations is significant. If the key leads to your organization's production and/or staging environments for software that you build or manage, for example, one stolen key could put potentially catastrophic events in motion for your organization.

While malicious actors can gain access in any number of ways, getting into your cloud environment via API is the equivalent of finding a wide-open door and walking in unchecked. Whoever holds that key will have direct access to what are likely your most valuable digital assets. And with API access, user-level security controls like multi-factor authentication become all but irrelevant. It's safe to say that exposure of your AWS API keys is a serious problem.

97%

of enterprise leaders consider a well-executed API strategy critical in driving their organization's growth & protecting revenue streams, yet according to a recent study, 84% of security professionals reported API security incidents over the past year.

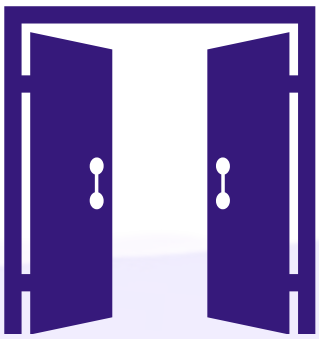

Security Risks of Compromised Secrets

Data Breaches and Information Leakage

One doesn't have to do much head scratching to infer the risks of leaked sensitive data. Nevertheless, when a company goes from booming enterprise to bankrupt within 12 months it's time to take notice. In October 2024, enterprise background check provider National Public Data declared bankruptcy, citing legal fallout from a massive data breach the company disclosed in August of the same year.

> *Experts estimate **899 million** unique SSNs of U.S. citizens were exposed by the hacker group "USDoD" after posting a database exfiltrated from National Public Data– a background check provider–on a dark marketplace website.*
> ***The company declared bankruptcy within 2.5 months.***

With more than 20 states adding civil penalties to the mountain of lawsuits, financial liabilities seem to have buried the organization.

Manipulation of System Configurations

Once inside your systems, attackers can change security settings, removing layers of security you have in place to guard against attacks. This allows the bad actors to set up *back doors* (connections to external locations for future data exfiltration), establish *persistence*, and more. (Persistence is established when a threat actor deploys malware that quietly reestablishes malicious connections if security teams ever find and close their back door.) Whatever their goal, a bad actor with access to your critical systems is a severe threat.

Control Over Virtual Servers

If attackers gain control even a single virtual server, they can do a lot of damage in a hurry. In a straight-forward attack, they might exfiltrate valuable intellectual property or other sensitive data directly from the server. However, they can also use that virtual server as a command center to launch ransomware attacks on other servers or undertake any number of suspicious activities.

Exposed SaaS API Keys

If an exposed API key grants access to a SaaS application that is only used in one part of your environment, the damage may not be as immediate, but it will be limited only by the creativity of the attacker. Compromised SaaS API credentials still enable unauthorized system access, potentially allowing malicious actors to create fraudulent accounts, harvest data, send spam, or even orchestrate DDoS attacks.

Lifted Encryption Keys

Think of it this way, you only encrypt data that's extremely valuable and could result in harm if stolen or exposed. If an encryption key is stolen, the primary risk is a complete data breach, where attackers can easily access and decrypt sensitive information like financial details, personal data, or trade secrets, leading to identity theft, fraud, reputation damage, legal repercussions, and significant financial loss. Further, exposure of regulated data could result in compliance violations, fines, and reputational damage.

Stolen Passwords

If a corporate account password–to any part of your environment–is leaked, the primary risk is unauthorized access to sensitive company systems and data, which can lead to significant financial loss, reputational damage, data theft, legal issues, and loss of customer trust due to privacy and security concerns. Malicious actors may also exploit a compromised password to gain access to critical corporate information and carry out any number of harmful other activities. For example, an email account password can allow them to use the compromised account for business email compromise attacks, malware attacks, and even lateral movement via social engineering.