

EBOOK

Credential sprawl: How AI increases the risks



Introduction

In Ancient Rome, the military had a daily “watchword” that soldiers used to enter the camp.¹ An official would inscribe the watchword on clay tablets, which were distributed throughout the various military units; if a tablet wasn’t returned, they swiftly tracked it down and punished the soldier who had failed to return it.

Clearly, one thing has been true from Ancient Roman times until now: if you want to stay secure, you need to know where your passwords are.

Unfortunately, keeping track of credentials is more difficult for a modern organization. Today’s companies have to manage an ever-growing number of credentials that go well beyond traditional passwords, such as developer secrets, passkeys, shared logins, API keys, SSH keys, service accounts, and SSO access tokens.

These credentials don’t live in one place. They live in browsers, scripts, developer environments, Slack messages, AI tools, config files, and sometimes plaintext spreadsheets.

The proliferation of credentials outside centralized visibility and control is known as “credential sprawl,” and attackers are eager to take advantage of it.² This problem is especially urgent due to the rise of AI-based tools and agents, which have not only increased the scale and scope of unmanaged credentials, but also present access and identity management challenges that tools like SSO and PAM aren’t equipped to handle.

¹ [Polybius The Histories](#), Loeb Classical Library, 1923

² [Too Many Secrets: Attackers Pounce on Sensitive Data Sprawl](#), Dark Reading, 2025

What causes credential and secrets sprawl?

Credential risks are hardly a new issue. However, in recent years, managing where and how credentials are used has evolved from a Herculean task to a Sisyphean one. That is to say: it was never easy, but at some point it became close to impossible. To understand why, we'll begin with an overview of some of the essential factors contributing to credential sprawl.

Poor password practices

The dangers of unmanaged passwords and poor password practices have been well known for years, and yet passwords remain one of the most common vectors for data breaches and hacks.³

1Password's 2025 annual report, [The Access-Trust Gap](#), found that two-thirds of employees admit to engaging in poor password practices, including:⁴

- Using the same passwords across multiple work accounts
- Never changing IT-default passwords
- Using the same password for both work and personal accounts
- Texting, emailing, or otherwise messaging passwords to yourself or a colleague

These poor practices can have devastating effects, like when attackers used credentials stolen in previous breaches to log into developers' repositories on GitHub, Bitbucket, and Gitlab. As journalist and analyst Robert Lemos reported, "The attack highlights the dangers of mishandling passwords... None of the services hosting affected developers' repositories found signs of a compromise. Instead, attackers logged onto them from an unrecognized Internet address using valid credentials and then deleted the victim's code."⁵

³ [Data Breach Investigations Report](#), Verizon, 2025

⁴ [Annual Report: The Access-Trust Gap](#), 1Password, 2025

⁵ [Password Reuse Misconfiguration Blamed for Repository Compromises](#), Dark Reading, 2019

“

The average company has a third of its apps outside SSO's protection.

Unmanaged SaaS and AI

As security analyst Francis Odum put it,⁶ “As organizations increasingly adopted SaaS applications, the need for enterprise-grade password management became more pronounced. Employees frequently relied on personal credentials for work accounts, increasing the risk of credential reuse and security incidents. While Single Sign-On (SSO) and Multi-Factor Authentication (MFA) became standard controls, they often failed to cover the full range of enterprise applications, leaving visibility gaps...”

SaaS sprawl has been a known problem for years in the cybersecurity industry.⁷ The average company manages over a hundred applications within their Single-Sign On (SSO) platform alone,⁸ but many apps are outside the reach of SSO.

In fact, 1Password found that the average company has a third of its apps outside SSO's protection. Our report also noted that, “One major indicator of how SSO is falling short is the amount of access that comes from employees whom IT believed to have been successfully offboarded. Over one-third (38%) of employees have successfully accessed a prior employer's account, data, or applications after leaving the company.”⁹

Now, AI is accelerating SaaS sprawl even further beyond what SSO was built for. One in four employees has used AI applications that weren't approved by their company, and over a third of employees admit to having knowingly disregarded their company's AI policies.¹⁰

Employees are experimenting with AI coding tools, browser extensions, writing assistants, data analysis tools, and agent platforms, often before IT has evaluated or approved them. Many of these tools don't integrate cleanly with enterprise SSO, and even when they do, adoption frequently begins outside official onboarding processes. Shadow AI poses serious risks since even innocuous apps can contain serious security flaws that could expose company data and credentials.¹¹

AI agents access credentials differently as well. They don't just require user logins; they often require API keys, OAuth tokens, service accounts, and other machine credentials to function. Those credentials may be stored locally, embedded in scripts, saved in browsers, or shared informally between teammates, far beyond the visibility of traditional identity systems.

Each unmanaged app and AI tool represents at least one unmanaged credential that an organization can't secure. The result is an ever-expanding layer of applications and credentials that exist outside centralized governance.

For a more in-depth look at SSO's limitations, read our ebook, [Why SSO is not enough for identity security](#).

⁶ [Inside 1Password's Enterprise Identity Transformation](#), Software Analyst Cyber Research, 2026

⁷ [Taking Control of SaaS Sprawl](#), 1Password, 2025

⁸ [Businesses at Work](#), Okta, 2025

⁹ [Annual Report: The Access-Trust Gap](#), 1Password, 2025

¹⁰ [Annual Report: The Access-Trust Gap](#), 1Password, 2025

¹¹ [Millions of Android AI Apps Hide a Dangerous Secret](#), Tech Times, 2026



Developer secrets

Developer secrets, such as SSH keys, API keys, service account credentials, environment variables, and server tokens, are the keys to an organization's most critical systems. Unlike user passwords, these credentials often operate quietly in the background, powering applications, infrastructure, automation, and now AI agents.

However, these secrets rarely live inside traditional identity systems. Instead, they're stored in code repositories, local environment files, CI/CD pipelines, cloud consoles, scripts, and collaboration tools. In the absence of consistent governance and purpose-built tooling, these secrets proliferate quickly, and often go unmonitored.

GitGuardian's 2025 report, *The State of Secrets Sprawl*, shows how rapidly this problem is accelerating. "In 2024, we found 23,770,171 new hardcoded secrets added to public GitHub repositories. This figure represents a 25% surge in the total number of secrets from the previous year." As they put it, "secrets sprawl is steadily worsening over time."¹²

Secrets sprawl can spread in a number of ways, including when developers accidentally expose secrets in public-facing code. However, GitGuardian's report highlights a more basic concern: "[While] source code management tools have been the primary focus of secrets detection... secrets appear wherever teams collaborate, often in collaboration and project management tools like Slack, Jira, or Confluence."¹³

Plaintext secrets being sent through apps like Slack represents a dangerously lax approach to secrets hygiene. Unfortunately, cybercriminals are aware of this trend. Dark Reading reports that "... cybercriminals and nation-state actors alike are following a proven playbook and capitalizing on 'bad secret hygiene' to further their campaigns."¹⁴

AI is now accelerating this dynamic. As developers use AI copilots to generate code, spin up infrastructure, or automate workflows, machine credentials are created and reused at greater speed. Without centralized oversight, secrets multiply across repositories, prompts, pipelines, and agents. All of this is expanding the identity surface far beyond what traditional identity and access management (IAM) and privileged access management (PAM) systems were designed to govern.

¹² [The State of Secrets Sprawl](#),
¹³ GitGuardian, 2025

¹⁴ [Too Many Secrets: Attackers Pounce on Sensitive Data Sprawl](#), Dark Reading, 2025

How AI makes credential sprawl worse

“

The short version: agent gateways that act like OpenClaw are powerful because they have real access to your files, your tools, your browser, your terminals, and often a long-term ‘memory’ file that captures how you think and what you’re building. That combination is exactly what modern infostealers are designed to exploit.

–Jason Meller, Vice President and Security Strategist, 1Password

The rise of AI and agentic AI have dramatically increased the productivity of the modern workplace. However, they’ve also accelerated the rate of credential sprawl just as dramatically, and their risks warrant closer analysis.

Unmanaged AI agent access

AI agents represent an entirely new class of identities; they require varying levels of access, but operate in ways that are frequently invisible to security tools.

As The Hacker News put it, “AI agents don't operate in isolation. To function, they need access to data, systems, and resources. This highly privileged, often overlooked access happens through non-human identities: API keys, service accounts, OAuth tokens, and other machine credentials.”¹⁵

All non-human identities (NHIs) pose credential risk, and the way that AI agents use them has increased their sprawl drastically. Figures vary, but in 2025, there were somewhere between 82¹⁶ – but potentially up to 144¹⁷ – NHIs for every 1 human identity in the average enterprise environment. Regardless, that number is growing fast.

More concerning is the fact that many of these machine identities have highly privileged levels of access, often without the level of scrutiny that would typically be applied to highly privileged users. In fact, a recent study found that 1 in 20 NHIs carries full-admin privileges even though only 38% of total NHIs had been active within the last 9 months.¹⁸

¹⁵ [The Identities Behind AI Agents: A Deep Dive Into AI & NHI](#), The Hacker News, 2025

¹⁶ [Machine Identities Outnumber Humans by More Than 80 to 1](#), CyberArk, 2025

¹⁷ [The NHI & Secrets Risk Report](#), Entro Labs, 2025

¹⁸ [The NHI & Secrets Risk Report](#), Entro Labs, 2025



The proliferation of Agentic AI creates identity sprawl and a critical access gap. AI agents operate probabilistically, needing access to internal systems via credentials (passwords, API keys) often designed for human use. Because existing access solutions were not designed for dynamic, probabilistic machine identities, developers often resort to hardcoding secrets. This leads to over-privileged agents, limited auditability, and elevated data loss risk.

–JFrancis Odum, Inside 1Password’s Enterprise Identity Transformation²²

What this means is: AI agents are being given increasingly powerful levels of access, that access is often going unmanaged by security teams, and agents are often holding onto powerful privileges beyond the point where they need them.

Agentic applications and capabilities are evolving at unprecedented speed, and new tools are often being adopted before their risks are understood. Jason Meller, VP and Security Strategist at 1Password, wrote two [blog posts](#) on how powerful – and frightening – these tools can be.^{19,20}

While OpenClaw certainly garnered some attention, its issues aren’t isolated to one tool alone. In MIT’s “AI Agent Index,” researchers found that the majority of agent developers share little about their tool’s security. “25/30 agents disclose no internal safety results, and 23/30 agents have no third-party testing information.”²¹ Rather, OpenClaw is an indicator of how severe the security risks can be when AI agents are given unmanaged levels of access; its popularity, and its security risks, have quickly forced security teams to reckon with the fact that the standard enterprise perimeter is not equipped to handle the issues of agentic AI.

AI multiplies bad credential practices

AI-based tools are also exacerbating credential sprawl by replicating poor credential security practices.

Vibe coding (using generative AI to write code) tends to reproduce poor security habits. For example, one largely vibe-coded platform, Moltbook, was quickly found to have a misconfigured database within it that exposed over a million API authentication tokens, along with email addresses and private messages.²³

Again, this isn’t exclusive to a single platform. GitGuardian analyzed the use of Copilot – Microsoft’s AI assistant (used for vibe coding, among other things) – and they found that repositories with Copilot active are 40% more likely to have at least one leaked secret.²⁴

On the whole, vibe coding can also enable employees with less coding experience, and therefore less coding security training, to push through code that hasn’t received some of the standard checks and scrutiny that should be applied toward any code’s security.

¹⁹ [It’s incredible. It’s terrifying. It’s OpenClaw.](#), 1Password, 2026

²⁰ [From magic to malware: How OpenClaw’s agent skills become an attack surface](#), 1Password, 2026

²¹ [The 2025 AI Agent Index](#), MIT, 2026

²² [Inside 1Password’s Enterprise Identity Transformation](#), Software Analyst Cyber Research, 2026

²³ [Vibe-Coded Moltbook Exposes User Data API Keys and More](#), Infosecurity Magazine, 2026

²⁴ [The State of Secrets Sprawl](#), GitGuardian, 2025



Traditional identity security is falling behind

Monitoring how employees use and store credentials has always been challenging. But AI fundamentally changes the identity security model.

AI tools and agents don't authenticate, store, or use credentials the way humans do. They rely on embedded tokens, API keys, service accounts, and programmatic access patterns. They operate continuously, duplicate easily, and often persist long after their original purpose has ended.

Traditional identity security tools were designed for human behavior, with interactive logins, session-based authentication, and clearly defined privilege tiers. They were not designed to govern autonomous software identities that scale and authenticate programmatically without supervision.

In a way, this is almost by design. As Saumitra Das put it in an article for Corporate Compliance Insights, "By nature, autonomous agents are trained to find the easiest and most efficient way to complete the assigned job. This means that they can often identify ways around guardrails..."²⁵

Traditional access control methods are quickly proving to be inadequate, as AI and event-driven automation create NHIs at a scale we haven't seen before. As TechTarget reported, "Most legacy IAM and privileged access management (PAM) tools were never designed to handle that level of volume and churn."²⁶

The article goes on to point out some of the issues related to how NHIs use credentials, including:

- NHIs use a broad array of authentication methods, like JSON tokens, cloud IAM roles, OAuth2 secrets, and API keys. Each of these has its own unique security needs.
- NHIs are often given outsized access and long-lived credentials so that teams can ensure the tool will have the access needed to automate various business processes.
- Anomaly detection can't always notice when something has gone wrong with an AI agent, since they don't really have "normal" behavioral patterns to deviate from.

Each of these factors can seriously damage the efficacy of a company's security stack.

²⁵ [Decoding Duty of Care in the Agentic AI Era](#), Corporate Compliance Insights, 2026

²⁶ [CISO's guide to nonhuman identity security](#), TechTarget, 2026

The costs of credential sprawl

What happens when credential sprawl runs rampant through a company? The costs manifest in a variety of ways, from an increased blast radius in the event of a breach, to time-consuming manual processes to manage security posture, compliance, and incident response.

Compliance failures

IT and security teams are consistently faced with the difficult task of achieving and proving compliance with regulatory standards like SOC 2, PCI DSS, ISO 27001:2022, and HIPAA.

Each of these standards has requirements related to the secure use and storage of credentials. For instance, PCI DSS requires that, “Audit logs capture all changes to identification and authentication credentials...”²⁷

SOC 2 similarly has various requirements related to how companies provision access to credentials, including requirements dictating that “Your organization should implement processes to remove credential access when an individual no longer requires such access.”²⁸ It’s worth noting that SOC 2 extends these requirements not only to user credential access, but to how “internal and external infrastructure and software” access credentials.

Regulatory bodies, on the whole, expect companies to prove that they’ve done their due diligence to protect sensitive information. “Due diligence,” in the case of managing credentials, means implementing essential tools to give admins oversight over where and how credentials are being used. Credential sprawl fundamentally undermines a company’s ability to do so.

Furthermore, while security tools may be falling behind the AI boom, regulatory bodies aren’t likely to cut companies any slack. If anything, they’re increasing their scrutiny. As Itamar Apelblat pointed out in an article for BleepingComputer, “In each of these frameworks, the organization is accountable for what happens to regulated data and regulated workflows. When AI agents are the ones acting inside those systems, accountability doesn’t disappear.”²⁹

²⁷ [PCI DSS: v4.0.1](#), PCI Security Standards Council, 2024

²⁸ [SOC 2 Controls CC6: Logical and Physical Access Controls](#), Hicomply, 2025

²⁹ [AI Is Rewriting Compliance Controls and CISOs Must Take Notice](#), BleepingComputer, 2026



Risk exposure

It's not hard to understand why compliance standards place so much emphasis on credential and access management. Simply put: credential sprawl greatly increases an organization's risk of cyber attack.

Compromised credentials are the single most common entry point for attackers,³⁰ and have been for some time; 50% of CISOs who've experienced a material breach in the last three years identified compromised credentials as a root cause.³¹

Credential sprawl significantly increases a company's attack surface. Each credential that's stored without security and IT oversight presents an opportunity for bad actors to breach systems. And with credential sprawl spreading so rapidly, companies are facing more risk than ever.

In 2025, IBM reported that shadow AI accounted for 20% of breaches, and 97% of AI-related security breaches involved AI that didn't have proper access controls. IBM also points out, "... that data was most often stored across multiple environments, revealing just one unmonitored AI system can lead to widespread exposure."³²

Incident response

Breach remediation and incident response are already costly and time-consuming processes. Credential sprawl is only worsening these issues. Shadow AI, for instance, adds more complexity and cost to breach response; a breach involving shadow AI can cost up to \$670k more than a comparable breach that didn't involve it.³³

According to GitGuardian, 70% of secrets that were leaked in 2022 were still valid in 2025.³⁴ That's a deeply worrying figure, indicating that compromised credentials aren't being remediated by any standard business process; they're not expiring automatically or being rotated by teams. And since these existing issues aren't being remediated, with every new breach, IT and security teams are buried under a pile of risks that just keeps growing in scope and complexity.

As TechTarget reported, NHIs and agentic AI complicate this issue further: "Since many organizations use NHIs to link cloud environments... secrets are often duplicated or reused across multiple systems, making remediation and rotation difficult if a single identity is compromised."³⁵

³⁰ [Data Breach Investigations Report](#), Verizon, 2025

³¹ [Annual Report: The Access-Trust Gap](#), 1Password, 2025

³² [Cost of a Data Breach Report](#), IBM, 2025

³³ [Cost of a Data Breach Report](#), IBM, 2025

³⁴ [The State of Secrets Sprawl](#), GitGuardian, 2025

³⁵ [CISO's guide to nonhuman identity security](#), TechTarget, 2026

Solutions for credential sprawl

Traditional IAM strategies aren't always going to be enough to manage the issues of credential sprawl. Rather, teams will require a multi-pronged effort that seeks to approach the problem from several angles.

Credential management

The issue of credential sprawl starts with password management. Unfortunately, even though passwords have plagued security teams for years, many companies still don't have a solid handle on them – and today's teams run more tools with more logins than ever.

Without a clear strategy, credential sprawl spreads unmanaged. As we explored in more detail in our recent blog,³⁶ [credential management](#) requires a strategy that's built to address how credentials are being used. Teams need systems around:

- **Coverage**, meaning what credentials to protect: passwords, passkeys, API tokens, SSH keys, NHI, and AI agent secrets.
- **Control**, meaning how those credentials are managed: where they can be stored, how they're shared, what rules apply, and how those rules are enforced.
- **Lifecycle**, meaning how credentials change: creation, ownership, rotation, revocation, and proof, especially as roles and privileges change for human and machine identities.

The first step to meeting all of those requirements is an enterprise password manager (EPM). EPM covers more than passwords alone, and it's a critical element to each of the pillars of a strong credential management strategy. As the analyst and researcher, Francis Odum, reported, "1Password's architectural anchor is its [Enterprise Password Management \(EPM\)](#) core. This zero-knowledge vault serves as the singular 'system of record for all workforce credentials,' spanning both human users and non-human identities (NHI)..."³⁷

Despite this, 1Password found that only about 38% of employees use a company-provided password manager.³⁸

³⁶ [IAM stops at sign-in. Your credentials do not.](#), 1Password, 2026

³⁷ [Inside 1Password's Enterprise Identity Transformation](#), Software Analyst Cyber Research, 2026

³⁸ [Annual Report: The Access-Trust Gap](#), Password, 2025

“

This zero-knowledge vault serves as the singular ‘system of record for all workforce credentials,’ spanning both human users and non-human identities...”³⁷

An EPM like 1Password’s is a mission-critical tool for companies to rein in credential sprawl and manage Agentic AI use.³⁹ It centralizes visibility into how credentials are used, allowing admins to enforce principles of least privilege access through role-based vault access. Structured onboarding and onboarding workflows mean that users are only given access to the credentials, passkeys, and secrets that they need to do their jobs. And critically, EPM extends protection into developer workflows and AI-powered automation without introducing friction.

Since credentials are encrypted, teams can ensure that they can’t be accessed by infostealers and other targeted attacks. 1Password’s breach monitoring also informs users and admins as soon as possible if a managed credential has been compromised.⁴⁰

In short: the credentials to every workplace app stay secured and centralized where IT can easily oversee employee access, manage onboarding and offboarding, and measure the strength and security of their password ecosystem.

It’s worth noting an essential element of EPM’s efficacy: credential governance must be deployed organization-wide. Businesses have to enforce credential management for every person, agent, secret, and workflow. Companies cannot stay secure by only protecting part of the identity surface.

Secrets management for AI and business-led IT

Broadly speaking, credential sprawl often comes down to the push and pull between security and productivity. The rise of AI has placed this conflict in stark relief: employees, and developers in particular, adopt AI to improve productivity. They often see security tools as intrusive blockers to their improved workflows.

1Password doesn’t just improve [secrets management for developers](#), it removes friction. 1Password’s developer tools let teams securely vault secrets and make them available at runtime as developers code, so that they can work securely without interrupting workflows.

When it comes to agentic AI use, 1Password has also taken steps to let teams take advantage of the benefits of AI-assisted coding without ignoring the risks. Our Cursor integration “... gives developers a secure, just-in-time way to ensure required secrets are made available to Cursor’s AI agents via 1Password Environments. The result is an AI-native development workflow where... secure access becomes a natural part of writing and running code.”⁴¹

This is just the beginning. 1Password is building the foundations of runtime access governance for AI agents and machine workloads, to empower IT and security teams with capabilities to discover and manage how AI and AI agents are being used in their company – and how those systems are using credentials.⁴²

This is the next frontier of credential management: governing not just who logs in, but how software identities authenticate, operate, and persist across environments.

³⁹ [Securing identities starts with 1Password: Enterprise Credentials Management](#), 1Password, 2025

⁴⁰ [1Password Enterprise Password Manager product page](#), 1Password, 2025

⁴¹ [Cursor with 1Password](#), 1Password, 2025

⁴² [Inside 1Password’s Enterprise Identity Transformation](#), Software Analyst Cyber Research, 2026

SaaS management

Credential sprawl and SaaS sprawl are irrevocably intertwined. For IT and security teams to effectively determine where and how credentials are being stored, they need to know what applications their employees are using.

The unfortunate nature of SaaS sprawl, though, is that it's next to impossible for teams to find the time or manpower to take control of it manually.

1Password SaaS Manager solves this problem through automation. With over 40,000 app integrations, it lets teams build and maintain a complete inventory of the apps their employees use – including the apps that can't be secured behind SSO. That includes capabilities for continuous app discovery to illuminate the use of shadow IT – and shadow AI apps – across an organization.

With automated onboarding and offboarding workflows, teams can also ensure that employee access to apps is provided only when needed, without running the risk of unapproved access from improperly offboarded employees.

Identifying which applications are in use, whether they're company approved or not, is a critical step to making sure that every credential is being used and stored securely. A team cannot achieve comprehensive credential security if any part of their application surface is going unmanaged.

Do you know where your credentials are?

Credential sprawl is far from a new problem. But rather than improving, it only seems to be getting worse; multiple studies show that credential management practices are in a steady decline, as teams are faced with an ever-growing number of credentials across an ever-growing number of endpoints and apps. Credentials are hidden in codebases, Slack messages, AI chatbots, spreadsheets – and they probably still find a home on a sticky note or two.

Credential management has never been more crucial. In blunt terms: every unmanaged credential puts your ecosystem at risk. If credentials aren't being secured organization-wide, then your business can have untold numbers of unsecured access points.

Enterprise password management has never been an optional solution for companies that prioritize security at every level, but the impetus on rolling out fundamental credential controls has only become more pressing with the rapid rise of AI. 1Password is the critical solution for companies to reign in and control how credentials are used across their ecosystems. By building on the strong security of our password manager, we're building systems that will let teams manage credentials wherever they may be, from the spreadsheet to the AI agent.

There's never been a better time to start managing credential sprawl. [Reach out for a demo.](#)

Conclusion

1Password

We've seen that employees will find ways to work around security systems that interrupt their workflows, unintentionally exposing their company to risk. It's up to IT teams to adapt to how and where their team is getting work done. Productivity and security don't have to be a "one or the other" option.

That's why over 180,000 businesses rely on 1Password's Enterprise Password Manager to secure their business. With 1Password, companies are able to secure every sign-in, regardless of where your workforce is. 1Password is built with your team in mind, ensuring they're secure without slowing them down.

[Learn more about 1Password for credential management.](#)