

The CISO's Guide to Securing Data in the AI-ready enterprise

The four forces breaking traditional data security in the AI-first world and how security teams can stay ahead.



Custom content for Palo Alto Networks from Studio by Informa TechTarget



Introduction

The modern enterprise is being rebuilt around hybrid work, cloud, SaaS, and AI, making it increasingly difficult for organizations to keep their sensitive data secure. Data, users, apps and devices are distributed across the cloud, home, branch and everywhere in between, with 85% of work now taking place in the browser. Generative AI adoption has exploded, resulting in exponentially more data flowing through LLM models, SaaS apps, and unmanaged endpoints in ways that are increasingly hard to monitor and control.

While AI-led innovation has become mission-critical, it requires a major rethink on how we approach data security. GenAI tools, AI agents and copilots are already embedded in everyday work, but many enterprises lack central governance, leading to a concerning rise of 'shadow AI'. "The vast majority of customers I talk to today are starting from a point of view that the AI genie is completely out of

the bottle," says Tim Davis, Director of Data Protection at Palo Alto Networks; "They're getting pressure from employees who are using or wanting to use GenAI tools, and they're getting pressure from their boards because they're hearing about the productivity gains."

In light of this proliferation of GenAI, the question becomes; 'how do we enable this productivity gain without compromising our data?' It's not just about keeping adversaries out, but keeping sensitive data under control wherever it lives and moves—often between sanctioned and unsanctioned apps and devices. To do that, organizations must apply comprehensive data protection across all business channels with precise, AI-powered data classification and unified, proactive management. That's what we'll cover in this playbook, explaining how to mitigate data security risks in the AI era—without getting in the way of user experience.



Four forces breaking traditional data security in an AI-first world

Traditional data security programs were built for a simpler reality: fewer channels, clearer boundaries, and data that's easier to identify and control. In an AI-first world, things become much more complex. Data doesn't just 'move'—it gets copied into prompts, summarized, and shared through SaaS-to-SaaS integrations and AI plugins. This new reality exposes four forces that have all but rendered perimeter-based data protection models insufficient.

1 | Data sprawl and unstructured content

Sensitive data is no longer concentrated in a handful of systems that security teams can realistically inventory and lock down. Instead, it's scattered across systems and locations, and continuously moves, with AI expanding data footprints with unstructured content that's hard to classify using traditional data loss prevention (DLP) approaches.

“Data formats are much richer since the advent of AI. We have multimodal data in the form of pictures, text, audio and more. Because of that, the process of identifying sensitive data across the enterprise becomes much more important. Traditional DLP solutions quickly hit their limits in that kind of environment.”

PRASIDH SRIKANTH

Senior Director of Product Management
Palo Alto Networks

Traditionally, data security was primarily optimized for sets of regular patterns or for structured data sets, so they struggle to keep up when sensitive data is part of unstructured content, such as AI prompts in natural language, or when such data moves across multiple locations.

2 | Insider and human-centric risk

In the AI-first world, every knowledge worker becomes a high-velocity data-mover. That means they become another source of risk—a single point of failure where sensitive data can end up beyond the control of the organization's policies. Such incidents are rarely a result of malicious intent, but well-meaning employees trying to work faster and more efficiently.

"The number-one concern around data security today is what information is being put into prompts and being sent to these GenAI applications," says Davis; "GenAI apps are very data-hungry, and the more data you give them, the better their responses will be. There's an incentive for people who are looking for those productivity gains to surrender lots of data to these tools."

The problem is that most users don't really think about how the data is being stored and used by the GenAI application. This makes it harder to distinguish between normal work and risky behavior, but addressing the problem requires more nuance than simply blocking access to AI tools.

3 | Data security versus user experience

Data security and user experience are often viewed as competing disciplines, with convoluted security tools encouraging employees to seek risky workarounds in order to maintain a good user experience. Rigid, context-unaware approaches, such as blocks, brittle policies and false positives, have long been synonymous with friction and, in a world where people expect to use AI, that friction can backfire.

"A lot of times, if you block things, you impact productivity," says Davis; "Then, when employees don't understand why something's blocked, they get frustrated. It's usually preferable to turn on alerting and notifications, instead of actually stopping users from doing something. It gives them a speed bump and tells them what the risks are, which they probably wouldn't be aware of otherwise."

Traditional DLP solutions often apply blanket actions, such as blocking users from sending sensitive files, regardless of their contents or context. Because they're not context-aware, they do little to educate users, resulting in a disruptive experience where teams are bombarded with false positives.



4 | Complexity and tool sprawl

In the AI-first enterprise, complexity isn't just about having too many tools or vendors, but about losing visibility and control as data flows through a rapidly expanding SaaS and GenAI ecosystem—including copilots, plugins and integrations that often live outside central governance. Without broad discovery, security teams can't easily determine which tools are receiving sensitive data or even which tools are in use in the first place.

"What we see in most organizations, when we start to give them visibility into their current AI usage, is that they often have dozens of different applications for solving just one business problem," says Davis; "That's a huge attack surface, where everybody is just picking the tool they like. You've got to be able to push that down to one or two apps that your security teams actually know about."

Compounding these challenges is the fact that even identity is no longer stable or even purely human. AI-augmented workflows are increasingly autonomous, with access and actions happening through service accounts, API tokens, plugins, copilots and other nonhuman identities.





Designing an AI-ready data security approach

These four forces demand a new response to data security, one that enables AI-led productivity while still allowing teams to discover where sensitive data lives, classify it with precision, and protect it with consistent controls wherever it moves. This requires treating every access path—browser, SaaS, AI tools, IaaS, web, email and endpoints—as equally important in order to apply consistent controls wherever data lives and moves.



Start with the data, not the perimeter

Legacy security models assume a clearly defined perimeter—a boundary between the enterprise's owned assets and the public internet. In the AI-first enterprise, that concept is largely obsolete, with increasingly vast amounts of data being routed through GenAI tools, plugins and APIs.

In this environment, the most reliable anchor is the data itself: what's sensitive, who can access it, where it's allowed to go, and which interactions are considered 'normal' versus high risk.

“Data is now at the edge of every interaction—browsers, SaaS applications and GenAI tools, to name a few. The first step is to have a governance strategy in place. Then, once you have that, you need identity—human and machine identities mapped to their specific use cases and the data they require.”

PRASIDH SRIKANTH

Senior Director of Product Management
Palo Alto Networks

Here's what that looks like in practice:

Define what 'sensitive' means to your enterprise

Create a classification model that reflects your business environment, taking into account assets like customer data, source code, M&A documents, regulated data and user credentials.

Map access in terms of identities instead of networks

In addition to employee and contractor identities, be sure to include machine identities, such as service accounts, API tokens, plugins and copilots, especially now that AI workflows are becoming increasingly autonomous.

Establish 'safe zones' and 'unsafe zones'

Decide where sensitive data can flow freely, such as between sanctioned collaboration tools, and where it requires additional controls, as might be the case with external AI tools and unmanaged devices.

The approach outlined above is nothing new, but it's doubly important in the AI era, where what really matters is what an identity is trying to do with sensitive data and through which channel.





Layer on AI to improve classification

With robust data discovery and classification in place, AI itself becomes useful in a very specific way: improving classification accuracy across highly variable, unstructured and fast-changing data flows. This approach also reduces alert noise, since AI-augmented classification brings contextual understanding to GenAI interactions.

“Traditional DLP is very content-driven,” says Davis; “It’s highly effective for looking at data content, identifying sensitive information, and stopping it from moving to places you don’t want it to move. But when we talk about AI use cases, you have to think about intent. When I’m using a GenAI tool, I’m prompting, asking questions, and trying to get to some outcome. So, we’ve entered into this realm of content plus behavior.”

Whereas traditional DLP requires security teams to explicitly define what needs to be protected—and blocked—AI-powered DLP introduces contextual and behavioral understanding. That means support for:

Broader coverage of data types and actions

AI allows DLP to work with a wider range of data types, such as multimodal inputs, file types and generated outputs.

Contextually aware classification

AI-powered DLP considers the channel and the user action, allowing it to distinguish, for example, between a file being shared internally and a prompt to an external GenAI tool.

Continuous feedback loops for tuning

Modern DLP can be a powerful accompaniment to user training, since it can provide real-time guidance on what to share or not share, and it can use feedback to reduce noise and focus policy enforcement on the highest-impact risk scenarios.

By classifying and controlling AI use within the right contexts, enterprises can start rolling out AI tools at scale without dramatically expanding their risk exposure.





Unify policy enforcement across channels

Having a single policy framework eases management complexity and applies consistent controls for easier, faster audits and cross-channel incident response. It's impractical to have a separate policy for every channel, and while 'unified' doesn't mean that every channel has to be identical, it does mean enforcement is consistent even when mechanics differ.

"I recommend defining one global data security policy based on business intent and applying it in a channel-appropriate way," says Srikanth; For example, with an email, you might look at the receiver and sender, whereas in an application like ChatGPT, you look at the prompts or, if it's a file-sharing app, you look at the files being shared."

The most effective approach is a three-layer policy model:

Intent layer (enterprise-wide)

For example, your goal might be to prevent customer data from leaving sanctioned business channels.

Channel layer (implementation)

Each channel, such as prompts, file uploads and browser copy and paste, has its own technical conditions and constraints.

Response layer (outcomes)

Possible responses, depending on intent and channel, may include blocking, encrypting, quarantining or simply alerting/warning the end user.

When security teams have a single control framework, rather than an inconsistent patchwork, policy enforcement becomes more consistent across the enterprise.



Converge security and user experience

If a unified security policy is the ‘what’, then user experience is the ‘how’. In the AI-ready enterprise, security controls have to meet people where work actually happens—mostly inside browsers and SaaS or GenAI apps—without turning every interaction into a stop-start workflow.

“Once you have a unified architecture, you can establish a common understanding of what an incident looks like and what constitutes risky behavior,” says Srikanth; “That way, you have the same level of telemetry, response actions and coaching experiences.” This matters because user experience and security outcomes have become inseparable—when you have data security controls delivered across channels consistently, users get predictable, low-friction guardrails that don’t keep getting in the way of their work.

Here’s what that convergence looks like in practice:

Default to coaching for low- and medium-risk actions

Instead of simply blocking everything, use warnings and justifications and nudge users into safer alternatives for lower-risk actions.

Reserve hard enforcement for high-impact scenarios

Only block regulated data types, repeat risky behavior and unsanctioned destinations where there’s a clear and unacceptable risk.

Treat performance as a security feature

The less friction people feel, whether it’s prompting AI, sending an email, or anything else, the less incentive they have to try to bypass controls.

With controls that protect sensitive data without becoming a bottleneck that employees attempt to work around, security and user experience stop being competing disciplines and work together.



Prove business value to stakeholders

An AI-ready data security program only scales if it can be justified in business terms—in other words, outcomes stakeholders care about, such as reduced risk exposure, lower infrastructure and licensing costs, fewer support tickets, and enhanced employee productivity.

“When we’re talking about GenAI, and SaaS security in general, the most important thing is reducing the amount of shadow IT and shadow AI,” says Davis; “When you do that, you’re reducing the attack surface. You can push it down to just a handful of apps that you can monitor and understand. That’s the first and easiest win, because it reduces the stress of false positives and improves productivity. If you see a 95% reduction in false positives, as our technology is doing, that has a material impact on how the business operates in terms of staffing and resource allocation.”

When seeking buy-in for their data security initiatives, there are several business-friendly metrics CIOs can refer to:

Risk reduction

With fewer exposure paths for sensitive data, businesses have fewer costly, high-severity incidents to worry about, and shadow AI becomes much more manageable.

Operational efficiency

With fewer false positives to chase after, investigation workloads get lower, and triaging incidents requires less staffing and other resources.

Employee productivity

Fewer unnecessary blocks and support tickets translate into smoother adoption of AI tools throughout the enterprise.

With the right metrics in place, CIOs can show that AI-ready data security is more than just a defensive spend, but a strategic capability that forward-looking enterprises can build on.

What's next?

Building an AI-ready enterprise means accepting a new reality, one where AI is embedded in everything your employees do. That involves sensitive data moving through browsers, SaaS platforms, APIs, GenAI tools and more—often much faster and at a greater scale than security teams can track manually.

But the winning approach isn't to slow the business down by simply blocking AI tools. It's about putting data-centric guardrails in place that can discover sensitive data where it resides and moves, classify it with precision, and enforce policies consistently across every channel. That way, security stops being a daily productivity tax and instead drives innovation without adding risk.

Palo Alto Networks' Data Security makes this possible with precise AI-augmented data classification, granular control and comprehensive protection for all your enterprise data.

[Learn more about Palo Alto Networks Data Security solution](#)





As the global AI and cybersecurity leader, Palo Alto Networks (NASDAQ: PANW) is dedicated to protecting our digital way of life via continuous innovation. Trusted by more than 70,000 organizations worldwide, we provide comprehensive AI-powered security solutions across network, cloud, security operations and AI, enhanced by the expertise and threat intelligence of Unit 42. Our focus on platformization allows enterprises to streamline security at scale, ensuring protection fuels innovation.

[Explore more](#)



Expert led. Impact driven.

Studio is Informa Tech Target's global content studio offering brands an ROI rich tool kit: Deep industry expertise, first-party audience insights, an editorial approach to brand storytelling, and targeted distribution capabilities. Our trusted in-house content marketers help brands power insights-fueled content programs that nurture prospects and customers from discovery through to purchase, connecting brand to demand.

[Learn more](#)