White Paper

Securing Your Data in the AI Era

A Data Security Council Whitepaper

DATA SECURITY COUNCIL

Sponsored by Netskope

Table of Contents

INTRODUCTION: MEETING THE CHALLENGE OF AI READINESS

AI is inescapable in today's technology sector. From industry events and investor calls, to boardroom reviews and all-hands meetings, AI is the topic on everyone's lips.

Adoption is rapid: 60% of CIOs polled by Morgan Stanley expected to have a generative AI project in production by the end of 2025[1]. And at the very top of organizations, AI is expected to drive transformational disruption. According to a recent 2025 Gartner® survey, "74% of CEOs believe that AI is the technology that will most significantly impact their industries over the next three years."[2]

But as they race to become AI-ready, organizations need to adapt their security approach for this new reality. For CISOs and data leaders, the urgent challenge is to prepare their organization for AI's impact on their policies, processes, and other elements making up their risk posture.

That's not simply about security alone. It's about how to balance safeguards with innovation, so that organizations safely unlock the benefits of becoming more AI-powered.

In this Data Security Council whitepaper, we reflect the attitudes and insights from council members on the pressing issue of AI readiness. The insights in this whitepaper are based upon conversations about how they, their customers, and their partners are preparing for it, at a time when the pressure on and expectations of security leaders have never been higher.

Working without a blueprint

The explosion of new, publicly available AI tools and systems since ChatGPT launched in November 2022 is unprecedented in its scale and speed.

Seemingly overnight, a new set of capabilities were accessible to organizations and employees. This created challenges that CISOs and data governance leaders are still adjusting to. As Ilan Dar, CISO and SVP Technology at AutoFi, put it: "We're barely prepared for what's out there now, let alone for what's coming."

Added to that uncertainty is the ongoing and seemingly accelerating speed of AI development. The pace is remarkable even compared to that seen with other recent technological shifts such as the boom in cloud computing. The difference between the cloud age and the AI era is that cloud was adopted over a period of a decade or more, while equivalent transformations in AI have taken just a couple of years. New technical breakthroughs and more powerful tools come along, seemingly every day.

Meet the Data Security Council

This whitepaper is based on Data Security Council discussions, including specific collaboration with the following members:

Sonali Bhagwat
Senior Director of Data Governance, Adobe

Ilan Dar
CISO and SVP Technology, AutoFi

Arthur Hedge
President, Castle Ventures

Karen Lopez
Senior Architect and Data Governance Consultant, InfoAdvisors

Venkat Valleru
Principal Information Security and Compliance Engineer, Informatica

[1] https://www.linkedin.com/posts/siddhanttrivedi_ai-enterprise-cio-activity-7353479251322417152-E3YV/
[2] Gartner Webinar, Executive Leaders, Build AI Literacy to Ensure AI Business Value Realization, September 22, 2025 https://www.gartner.com/en/webinar/755958/1711783-executive-leaders-build-ai-literacy-to-ensure-ai-business-value-realization GARTNER is a trademark of Gartner, Inc. and/or its affiliates.

This rate of change leaves CISOs operating "blind," with no clear idea of what new developments are coming soon and how they will affect their security posture. They are operating without a blueprint for the future, which means practitioners feel like they are in "a catch-up game constantly," according to Venkat Valleru, Principal Information Security and Compliance Engineer at Informatica.

For many security leaders, the goal is to plan as best they can, creating scenarios that try to anticipate potential implications and relevant security approaches. While there are inevitably many unknowns, this at least provides a framework for thinking about the future.

Toward AI readiness

In the rest of this whitepaper we explore how data and security leaders, from CISOs to data governance directors, are addressing the challenges of AI readiness in their organizations. We cover five areas specifically:

- The current state of play in organizational AI adoption
- The risks and threats from AI that concern them most
- How organizations are seeking to limit those risks and find the right controls
- The ongoing evolution of AI
- Directions for future discussion

AI is now an integral part of many organizations' digital environments. This Data Security Council discussion shares experiences and insights that can help CISOs and their peers benchmark their own approach.

THE STATE OF PLAY IN AI ADOPTION

Industry investment in AI continues to grow at a rapid rate. According to analysts at IDC, enterprise AI spending worldwide rose to $241.8bn in 2025—and is predicted to rise beyond $867.3bn in 2029."[3]

Those numbers tell only part of the story, however (and not just because a significant percentage of AI in any given organization comes through free service access). The reality is that AI adoption is a spectrum and progress along it is highly uneven. Some pioneers are already embedding sophisticated, private AI tools into their workflows, while the vast majority of organizations are still experimenting with relatively simple use cases.

Today's typical AI adoption journey can be divided into five steps, reflecting both where most organizations are now (steps 1-3) and where they're headed (steps 4-5). Each step contains potential risks that need to be managed too. The five steps are:

1. **Experimenting with AI tools**, and managing the risks of shadow AI that come with it
2. **Leveraging embedded AI in SaaS platforms**, without enabling unapproved data sharing
3. **Managing stand-alone AI apps**, and stopping data misuse and leakage
4. **Building private AI apps**, while avoiding common application vulnerabilities
5. **Deploying autonomous agents**, without granting over-permissive access

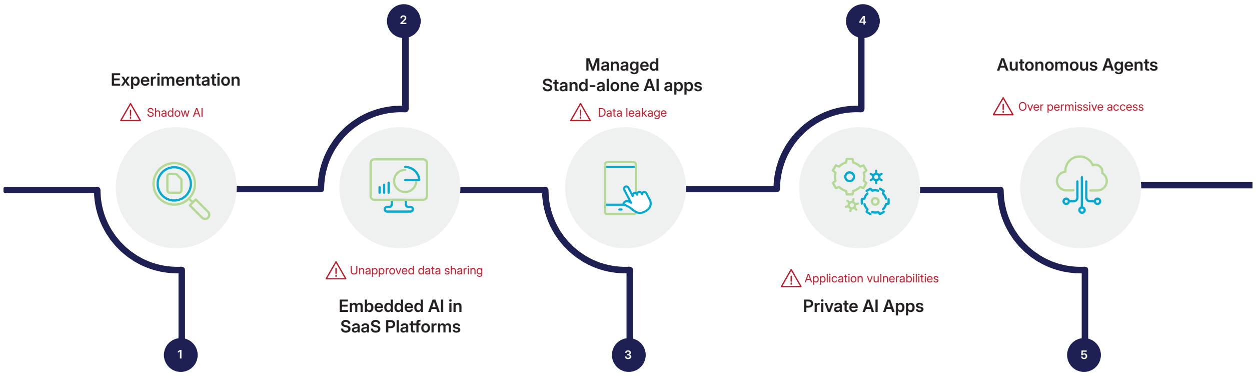

1
Experimentation
⚠ Shadow AI
2
⚠ Unapproved data sharing
Embedded AI in SaaS Platforms
3
Managed Stand-alone AI apps
⚠ Data leakage
4
⚠ Application vulnerabilities
Private AI Apps
5
Autonomous Agents
⚠ Over permissive access

Experimentation mode

Most Data Security Council members stated that they and the organizations they work with are still in the early stages of their AI evolution. In the experimentation phase, they are mainly making use of publicly available AI tools, especially large language models (LLMs) that can speed up admin work and help with brainstorming and research.

This is reflected in industry data too. Netskope's Threat Labs researchers recently reported that generative AI (genAI) apps are spreading rapidly across organizations. As many as 94% of organizations now use genAI apps, up from 81% a year before, with three times as many employees using them. A typical organization uses an average of 9.6 genAI apps now, up from 7.6 a year earlier[4].

At this experimentation stage, according to Arthur Hedge, President at Castle Ventures, "AI is used as a toolkit, but it's not embedded in workflows" and it doesn't involve rethinking how the business operates. This allows for relatively fast and low-cost gains in day-to-day work. Karen Lopez, Senior Architect and Data Governance Consultant at InfoAdvisors, reports as an example that some organizations are implementing applications such as customer service chatbots on their websites, viewing these as comparatively "easy" ways to show off AI capabilities.

[3] IDC, Worldwide Artificial Intelligence IT Spending Forecast, 2025–2029, IDC #US53688725, August 2025
[4] https://www.netskope.com/resources/threat-labs-reports/threat-labs-report-europe-2025

Shadow AI risks

The inevitable challenge with the experimentation stage of AI adoption is the risk of shadow AI. Without a chosen enterprise-wide tool in place, employees in effect bring their own LLM to work. And because they use personal instances of these models, it's often impossible for security teams to install behavioral guardrails, and as a result, they have no visibility to what information is being divulged or by whom.

Even in organizations not officially implementing AI projects, InfoAdvisors' Karen Lopez says, there is "a lot of shadow AI" going on. Free tools such as external notetakers that summarize video calls, for example, are frequently embraced without any consideration of data security.

The emphasis at many organizations today seems to be on embracing AI experimentation rather than protecting against security risks. There is a common view among Data Security Council members that this tone is often set from the top of the business. As AutoFi's Ilan Dar explains: "All [that] CEOs are hearing about is AI—and how it will allow you to cut costs and speed up delivery—so the demand on us has been significant in terms of '10x your output and security be damned.'"

The belief that "innovation outweighs the risk that comes with it" (in the words of Sonali Bhagwat, Senior Director of Data Governance at Adobe) appears to be highly prevalent.

However, Data Security Council members also report seeing the risks that AI brings, opening the C-suite's eyes to the importance of security. In the words of Castle Ventures' Arthur Hedge: "As someone who's worked in data security for 20 years, this is the most interested that business people have ever been in data security." Because these tools are so easy to use—and from a security perspective, to misuse—more company leaders are paying attention to the risks that come with them.

Growing sophistication

Among Data Security Council members, a distinction is generally made between AI used for internal activities such as marketing, sales, and engineering—which can be done through existing consumer-grade AI systems—and the more sophisticated work of building AI into the company's customer-facing products. AutoFi's Ilan Dar points out that the former approach of internal experimentation helps accelerate the organization's velocity and productivity, while rollout of the latter is a far bigger challenge. In his words, it can feel both "scary and fun" as the organization develops.

> "Many security leaders are often being forced to accept some level of risk and data leakage for the sake of productivity benefits, leading to a constant double-edged sword scenario that can be full of vulnerabilities. It's a very challenging balancing act."
>
> Ilan Dar, Chief Information Security Officer and Senior Vice President Technology, AutoFi

There's no denying that the bulk of organizations are in the early stages of their AI journey, but there is an active minority operating at the leading edge. Castle Ventures' Arthur Hedge reports that some of his customers in the financial services industry are "already building private AI apps to analyze investment data," and in fact he has "one customer with 100 AI pilots underway."

This reflects the reality of the AI adoption process. Most organizations are taking small steps. But for the outliers, AI tools can be transformational, and rapid implementation can make all the difference in business success. For security professionals, it means their remits are situational and will vary based on sector, mindset, technological sophistication, and risk appetite.

ADDRESSING KEY RISKS

> "When it comes to organizational usage of generative AI apps, I would say data leakage is probably the scariest thing for me. What information are we feeding into these systems? What are we allowing them to read through that could lead to a data leakage and oversharing challenge for us?"
>
> Ilan Dar, CISO and SVP Technology, AutoFi

Compared to a year ago, more organizations are using genAI, and more employees are using genAI. That expands the attack surface for each organization and creates ongoing headaches for security leaders.

But where do Data Security Council members see the greatest threats coming from? And what's worrying them the most about AI's impact? They call out six different risks.

Even before getting into cybersecurity issues specifically, there are foundational challenges that underpin how AI operates. One of these is the risk of **privacy issues** connected to how data is originally gathered and then used to power AI models and applications. Karen Lopez from InfoAdvisors captures some of the important considerations: "Are we using this data in the way we collected consent for? And are we potentially incurring a large liability by using it?" This raises the important issue of whether employees, customers, or partners are aware of how their data is gathered, shared, or repurposed for model training.

Likewise, poor **data quality** can undermine even the most advanced AI system. Missing, inaccurate, outdated, or biased data sets are common in many organizations. This can seriously hold back efforts to leverage AI tools or build proprietary ones. Problems with inconsistent data governance and management can often stretch back years, meaning they're difficult to solve quickly.

Lopez captures the concern: "Do people understand the technical debt in our source data?"

When scrutinizing the way in which AI tools are used, **data leakage** is generally considered the primary risk, standing out as the top threat discussed by Data Security Council members. Research from Verizon shows that around one in seven employees (15%) use generative AI tools on corporate devices, and most of them (72%) do so with personal email accounts[5]. This raises serious concerns as staff can put company intellectual property (IP) into these models, or personally identifiable information (PII) relating to customers. Furthermore, there is the risk of bad actors breaching an AI tool and exfiltrating company data. Council members generally believe that some level of data loss or leakage is inevitable when people rely on AI. The real question is how much exposure an organization is willing to accept in exchange for potential productivity benefits.

This also leads security professionals to worry about how their company's data is being used in **training third-party models**. Some employees are using new "free" AI features inside existing SaaS products without realizing that the vendor may use their data as part of their training process. As a result, some companies have had to retrain staff, update policies, and add safeguards to protect data privacy around AI.

Amid the AI hype, it can be easy to focus on the security risks that come with data and the prompts that users put into LLMs. But there's also a need to **secure AI artifacts** such as vector embeddings, which are still potentially valuable assets and require safeguarding. Vector embeddings—a compact way of turning text and images into numbers so an AI model can understand them—capture the original data's meaning, and can still reveal sensitive information or allow someone to reconstruct aspects of it. One way to trace and govern these vector embeddings is to apply data lineage principles to them, going beyond tracking the raw data inputs into a model to oversee other key artifacts.

[5] https://www.verizon.com/business/resources/infographics/2025-dbir-smb-snapshot.pdf

Lastly, as LLMs get more sophisticated, risks from **Model Context Protocol** (MCP) become more significant. MCP connects AI models and agents to external tools and data. While highly convenient in creating smarter LLMs, it also widens the AI system's access surface and thus its risk potential. As AI systems become increasingly interconnected in the years ahead, governing MCP links will be a vital part of enterprise security practices.

FINDING THE RIGHT CONTROLS

AI systems are becoming deeply woven into organizations, and with that comes new kinds of exposure. When models can pull in sensitive data, generate code, or make decisions at scale, even small gaps in oversight can lead to big risks. Applying clear controls is about making sure the systems that organizations rely on behave safely, protect information, and deliver value without creating unintended and harmful consequences.

In Data Security Council discussions, members talk about several different approaches to limiting the risks that come with AI deployment and finding a stronger sense of control. Some of these are being applied right now, while others are in the exploration phase as security professionals look to the future.

Starting with data

While it's essential that organizations control the use of AI systems, there's arguably a preliminary issue they need to consider first. Most organizations simply don't know the extent of the sensitive data they have or where it lives. As InfoAdvisors' Karen Lopez remarks: "You can't protect something you don't know about."

For that reason, the first building block of effective control comes in two parts: **data discovery and classification**. By locating, tagging, and labelling their data properly, companies are better placed to govern it well. Castle Ventures' Arthur Hedge highlights that many organizations store extremely important information—such as tax returns, passports, and even birth certificates—making careful classification a vital step.

Network-level approaches

One "brute force" method for establishing control over AI is simply to block access via **firewalls**. Security teams could put a rule in place to deny access to certain URLs, for example, of public AI chatbots on corporate networks. This blanket ban was a more common approach a year or two ago, but has become harder to maintain more recently, especially as attitudes toward AI become more favorable among the C-suite.

Similarly, many organizations have looked at the possibility of installing **browser extensions** to monitor AI use and block specific actions (such as uploading source code). This involves using a cloud access security broker (CASB) to block confidential information being entered into an LLM, or an endpoint protection tool on devices to detect unauthorized behavior.

Data Security Council members are also deploying **data loss prevention** (DLP) tools, with firm guardrails on what users can and can't do within AI systems. There are challenges with implementing these policies consistently given that, as AutoFi's Ilan Dar points out, "AI tools are so easily accessible on the web, and users can type in data, they don't even have to upload a file."

Access controls for LLMs and AI agents

That challenge is leading other security professionals to look at LLMs themselves as the appropriate control point for risky user behavior. One option would be to deploy an **LLM gateway** as a way to block sensitive data going into the AI system. Adobe's Sonali Bhagwat reports: "We're looking into gateways as a point to prevent restricted information from getting out or being exfiltrated either by employees or adversaries." While this approach is in its early stages, it reflects some of the new thinking among data security professionals.

Another popular control lever is an organization-wide AI system that controls **both the LLM as well as the data** it accesses. Some popular tools, such as Google Gemini or Microsoft Copilot, offer this functionality. Gemini can access Google Drive documents but will limit its access only to those files that the individual user has permission for. Copilot can analyze documents and data on corporate networks in the same permission-bound way. This embeds strong controls by default and offers a degree of reassurance to security teams.

With AI agents, which are becoming increasingly common in organizations, the key is to give them the right level of permission so they can access only appropriate data. AI agents often connect and draw upon complex ecosystems of datasets and applications, making it hard for security professionals to have effective oversight. The risks of over-permissive access are high, so it's important that practitioners implement robust measures such as zero trust frameworks, with least-privilege access controls.

Organization size

Approaches to security controls can look very different depending on the size of the organization. **Smaller organizations** can have additional challenges. They can secure the cloud and the endpoint device, but they don't have enterprise networks that connect them, so they can't use that as a plank in their security model. AutoFi's Ilan Dar explains: "Shadow AI becomes exponentially more difficult to manage and track in small companies, because they don't operate an enterprise network that can be controlled."

At the other end of the scale, some **large organizations** have been experimenting with highly customized approaches to security. One Council member referenced an organization that was trying to protect its sensitive financial information by giving questions to the compliance team to input into an LLM, so they could review the results and determine whether they provided the output back to the employee or not. That cumbersome approach is unlikely to last, given how unscalable it is. But it's a sign of how far some organizations are willing to go to try to retain control over crucial data in the AI era.

> "Simple blocking isn't going to be effective. Driven by work stresses and the magical promise of productivity, developers are highly motivated to use work-arounds, such as circumventing controls by spinning up their own unmonitored virtual machines in public cloud environments to run open-source genAI code."
>
> Karen Lopez, Senior Project Manager and Architect and Data Governance Consultant at InfoAdvisors, Inc.

AI technology is still in its infancy and the field is evolving fast. Recently, genAI tools have been supplemented by agentic AI, in which bots talk to each other with no human intervention. These AI agents hold out the promise of major upgrades in enterprise velocity and efficiency, and are therefore attracting attention from C-suite leaders.

Analysts have forecast significant expansion in how AI agents are used over the next few years, as well as the risks that come with them:

- IDC predicts that by 2026, 40% of all job roles in Global 2000 companies will require working with AI agents.[6]
- Gartner® forecasts that, "by 2028, 25% of enterprise breaches will be traced back to AI agent abuse, from both external and malicious internal actors."[7]
- As a result, Gartner also predicts that "40% of CIOs will demand 'guardian agents' be available to autonomously track, oversee, or contain the results of AI agent actions by 2028."[8]

As AI continues to develop, these new technical capabilities give rise to additional security considerations. Data Security Council members reveal that many leaders in the field are already beginning to think through these implications—even if, at the moment, they have more questions than definitive answers.

Agentic AI's data challenge

In many respects, agentic AI exacerbates the challenges that security professionals are grappling with already. Because AI agents operate autonomously once initial rules are set, it becomes even more important that they are bound by robust policies and controls in the first place.

Data Security Council members point to data guardrails as being particularly important. As Adobe's Sonali Bhagwat explains: "We have to classify our data and label it accordingly, so that agents know exactly which data they're allowed to use and train on."

> "Large companies are choosing to bring things in house, rather than attempt to fix the complex and pervasive security problems associated with external SaaS use. By hosting the models and data locally, they can guarantee the data never leaves their environment."
>
> — Arthur Hedge, President, Castle Ventures

This also extends to questions of data sovereignty. For organizations who operate in multiple jurisdictions and need to comply with different regulatory regimes, clarity on where data is collected, stored, and used becomes essential. AI agents could easily contradict data residency rules if their remits are not established carefully.

The rise of on-premises AI

The majority of enterprises today rely on third-party AI tools, either from specialist providers like OpenAI and Anthropic, or as part of existing SaaS applications from strategic vendors such as Google, Microsoft, and Salesforce.

However, the security risks that come with relying on external suppliers are leading some organizations (especially larger ones) to consider handling more of their AI development themselves. In the words of Arthur Hedge from Castle Ventures: "Rather than fix the security problems, and have to put in place a huge data governance solution, why not just bring it in-house?"

[6] IDC Press Release, IDC FutureScape 2026 Predictions Reveal the Rise of Agentic AI and a Turning Point in Enterprise Transformation, 23 October, 2025
[7] Gartner Press Release, Gartner Unveils Top Predictions for IT Organizations and Users in 2025 and Beyond, October 22, 2024 https://www.gartner.com/en/newsroom/press-releases/2024-10-22-gartner-unveils-top-predictions-for-it-organizations-and-users-in-2025-and-beyond GARTNER is a trademark of Gartner, Inc. and/or its affiliates.
[8] Ibid

While this isn't a realistic option for many, for the very biggest firms—with the budgets to invest in their own tools, and the IT skills to match—it could be a plausible option. Even at somewhat smaller firms, it could be possible for back-office IT teams to support key parts of the AI development process on-premises, without sending data out to third parties.

> "Over time, large public AI tools may see too many suspicious or manipulated results in their learning datasets. To maintain data integrity and ensure the reliability of their AI outputs, organizations are building private models trained only on their clean, verified, proprietary data."
>
> — Arthur Hedge, President, Castle Ventures

DIRECTIONS FOR THE FUTURE

Every organization today is on a journey to adopt AI across its operations. For data and security leaders, ensuring their company is AI-ready is one of the defining challenges of our era.

This Data Security Council whitepaper aimed to capture the experiences and attitudes of leaders at the forefront of this field. But we're still early in AI development, and it's a conversation that will continue for years to come.

Based on what we heard from Council members, five themes stand out as being especially important today and likely to fuel further discussion:

1. **Data leakage is becoming the defining AI security challenge**. As staff use free AI models or embedded features across SaaS tools, sensitive information (such as customer PII and company IP) flows into these systems almost inevitably. This is forcing security teams to decide how much risk they can tolerate.
2. **Demand is growing for granular, data-centric** controls. Organizations realize they must know what sensitive data they have, and where it lives, before they can protect it. This is driving them to focus on deeper data classification, and protection of less obvious AI artifacts such as vector embeddings.
3. **A shift from SaaS AI to private on-prem AI**. Larger enterprises, especially in regulated sectors, are exploring the potential of on-prem hosting and private LLMs to avoid sending sensitive data to external providers. This also gives them the benefit of finer control over training and usage.
4. **Security frameworks are struggling with agentic AI and over-permissive access.** As autonomous agents start connecting tools and pulling in context dynamically, new points of failure are created. This accelerates demand for controls that can constrain agentic behavior and define access based on data sovereignty regulations.
5. **AI security is moving toward control point enforcement.** Shadow AI is an enormous challenge, leading organizations to explore centralizing governance through LLM gateways, firewalls, and controls at the device or browser level. Whether this policy-driven access approach catches on remains to be seen.

Charting the AI Path

As we've explored throughout this white paper, the key to successfully leveraging AI in the enterprise is balance—balancing innovation with security, transparency with functionality, and risk with reward. AI offers tremendous potential to transform business operations, but with that potential comes the responsibility to protect data, ensure compliance, and maintain ethical standards.

The insights shared by industry leaders during our workshop underline the importance of a thoughtful, strategic approach to AI adoption. Whether it's implementing data governance frameworks, carefully managing AI's integration into existing systems, or anticipating future regulatory requirements, the path forward requires caution and creativity.

But this is just the beginning of the conversation. AI's challenges and opportunities are evolving rapidly, and staying ahead will require ongoing dialogue and collaboration. We invite you to join us in our upcoming events, where we'll dive deeper into these topics, share new developments, and continue to explore how we can responsibly harness the power of AI. Be on the look-out for our upcoming content, webinars, workshops, and events designed to keep you at the forefront of data security innovation.

Ready to lead the conversation?
Join the Data Security Council. Gain exclusive insights, build your reputation as a thought leader, and network with elite industry leaders to influence the future of data protection.

[Join the Community]