ecosystm.

Shadow AI: Risks, Exposure & Governance

AUTHORED BY

Darian Bird
Principal Advisor, Ecosystm

Sash Mukherjee
VP Industry Insights, Ecosystm

MAY 2026

Table of Contents

Introduction

Not all AI inside the enterprise is visible.

Employees are using AI tools on their own to get work done faster and more easily, often prioritising convenience and speed over formal approvals or approved platforms.

This is where Shadow AI starts to emerge; when these tools are used outside formal security and governance controls. The intent is usually benign, but the pattern is familiar from earlier waves of Shadow IT, where ease of access and immediate value ran ahead of governance.

What is increasing the challenge now is how far this behaviour is extending beyond simple, chat-based use cases. It is showing up through APIs, internal scripts, and embedded business processes. That shifts Shadow AI from individual user behaviour to something built into systems, where data can be processed automatically and repeatedly without formal oversight.

This changes the risk profile quite significantly. Instead of isolated, human-speed data exposure, organisations are now dealing with machine-to-machine data flows that are harder to track, monitor, or audit. Data can move and combine across systems without clear ownership or visibility, making consistent governance much harder to enforce. Organisations need to rethink how security and control work when AI becomes embedded in day-to-day execution.

The Expanding Shadow AI Risk Surface

The rise of consumer-grade AI tools, especially GenAI productivity tools that are easy to access, intuitive, and often free to use, has significantly changed the risk landscape for enterprise technology leaders.

Unlike traditional enterprise software, these tools can be adopted instantly by employees without involving IT teams, which makes governance harder to enforce.

Governance models have not kept pace with these newer patterns of AI adoption; most organisations still rely on controls designed for core systems and approved applications. While visibility and enforcement tend to be stronger inside formal environments, they are weaker at the edges where new tools, workflows, and usage patterns are emerging.

FIGURE 1

Enterprise AI Governance: Where Control Exists Today

9%
No formal policies

21%
Basic policies exist, but enforcement is inconsistent

50%
Governance covers core systems only

12%
Structured governance with ownership, approvals & sanctioned tools

9%
Governance embedded in workflows, improving visibility of AI use across the enterprise

Source: Ecosystm, 2026

While governance frameworks evolve, technology leaders are concerned with the operational implications, particularly the challenge of maintaining visibility and control as AI becomes embedded in everyday work. What was possible to contain within defined systems is extending into employee-led experimentation, workflow-level integration, and tool adoption outside formal channels.

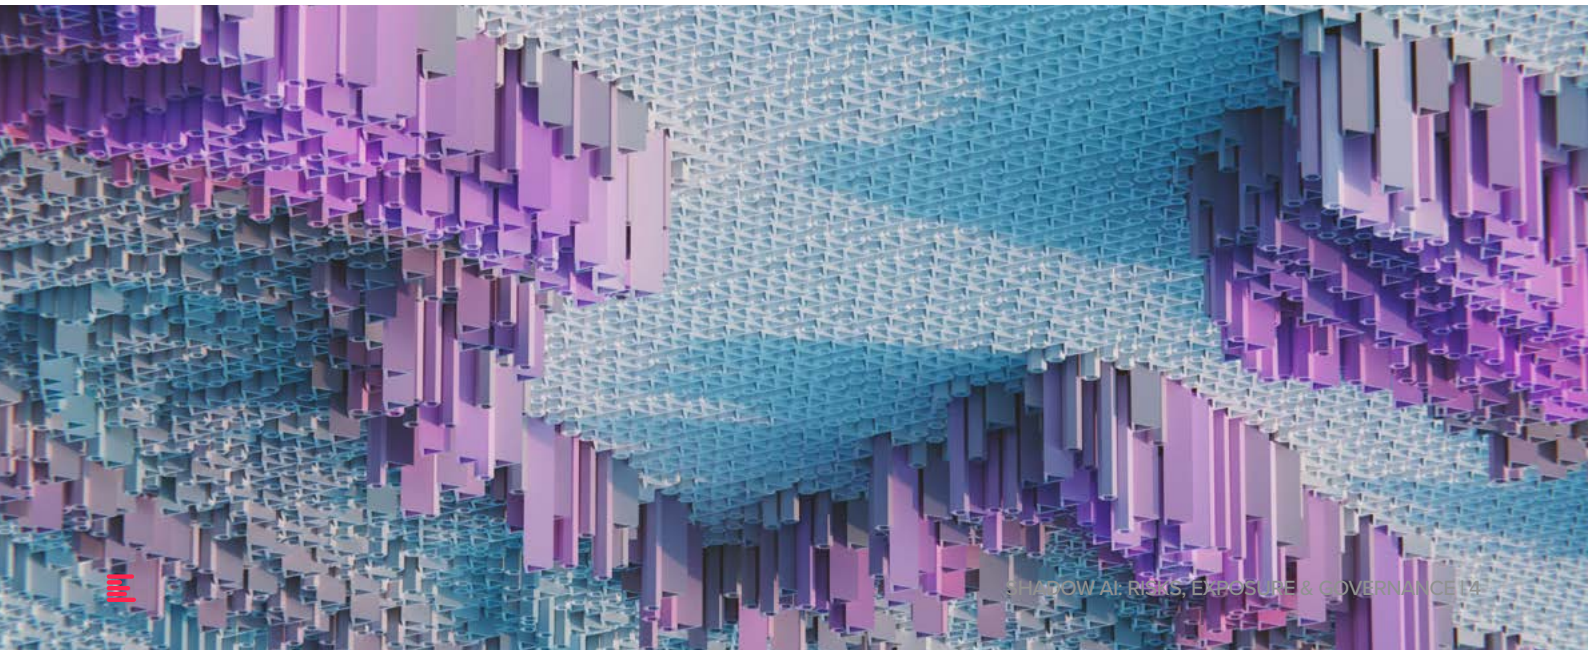

FIGURE 2

Biggest Challenges of Tech Leaders Today

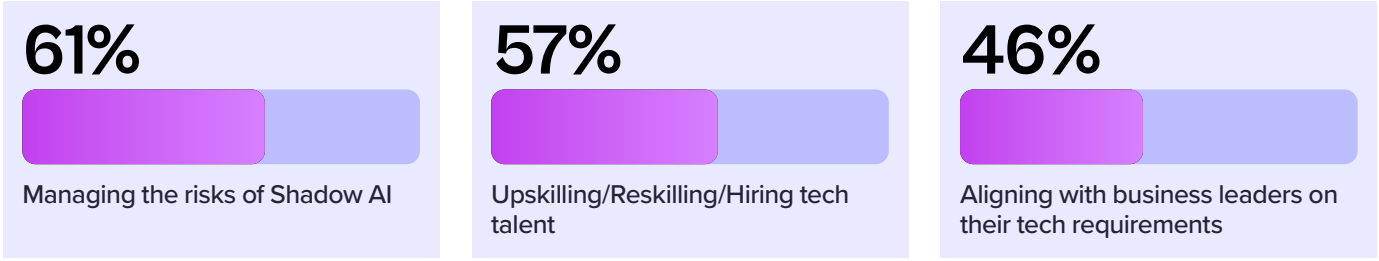

Source: Ecosystm, 2026

The risk has increased with greater familiarity with AI tools across the workforce, making experimentation more natural and more widespread. This risk of Shadow AI does not come from standalone chat-based tools today. It is emerging across multiple layers of the technology stack, where AI is embedded into browsers, SaaS platforms, developer tools, and automation systems, expanding both the surface area and complexity of unmanaged AI use.

Concern around Shadow AI has risen sharply – from 33% in 2025 to 61% in 2026.

Source: Ecosystm, 2026

Understanding Enterprise Exposure to Shadow AI

Data Leakage is the most immediate risk for organisations

Employees can paste sensitive information, including PII, internal documents, and intellectual property, into external AI tools that sit outside enterprise control. In many cases, there is limited clarity on how this data is stored, reused, or retained, increasing exposure to confidentiality breaches, compliance issues, and prolonged retention of sensitive information beyond organisational control.

Samsung Electronics banned the use of GenAI tools by its employees after it was discovered that a software engineer had uploaded sensitive internal source code to ChatGPT.

Regulatory & Compliance Risk emerges when AI tools are used outside approved processes and controls

This can result in breaches of privacy and data protection requirements, particularly where sensitive data is processed or moved across jurisdictions. Limited visibility into how AI is used also makes it harder to demonstrate compliance or trace how data has moved across systems, when required.

Data Provenance & IP Risk relates to uncertainty over where AI outputs come from and what they contain

Outputs may be influenced by copyrighted or restricted material, creating downstream risk if used in enterprise-facing work. Ownership is not always clear, which raises questions around reuse, attribution, and potential infringement.

Sovereignty & Foreign Exposure Risk is becoming more visible as organisations rely on models built and hosted outside their jurisdiction

With limited availability of frontier models, most enterprises have little choice over underlying infrastructure, increasing exposure to external legal frameworks and opaque data handling practices. This becomes more sensitive when strategic or commercially important data is involved.

Employees at the Pentagon allegedly used Chinese LLM, DeepSeek for several days until the Defense Information Systems Agency blocked access to the service. DeepSeek's privacy policy explicitly states that data is stored on servers located in China and governed by Chinese law.

Autonomous & Agentic Risk is emerging as AI systems move from generating outputs to executing actions

Without clear guardrails, systems can trigger workflows, move data, or interact with other systems in ways that were not really intended. These risks are harder to detect because they often only surface after actions have already taken place.

Model Reliability & Decision Risk shows up when outputs are accepted at face value despite being incorrect or biased

Even when responses sound credible, they can contain errors that influence downstream decisions.

Legal analyst, Damien Charlotin, has documented over 1,400 instances of AI hallucinations (till date) in legal proceedings globally, where filings included fabricated case law or incorrect references generated by AI.

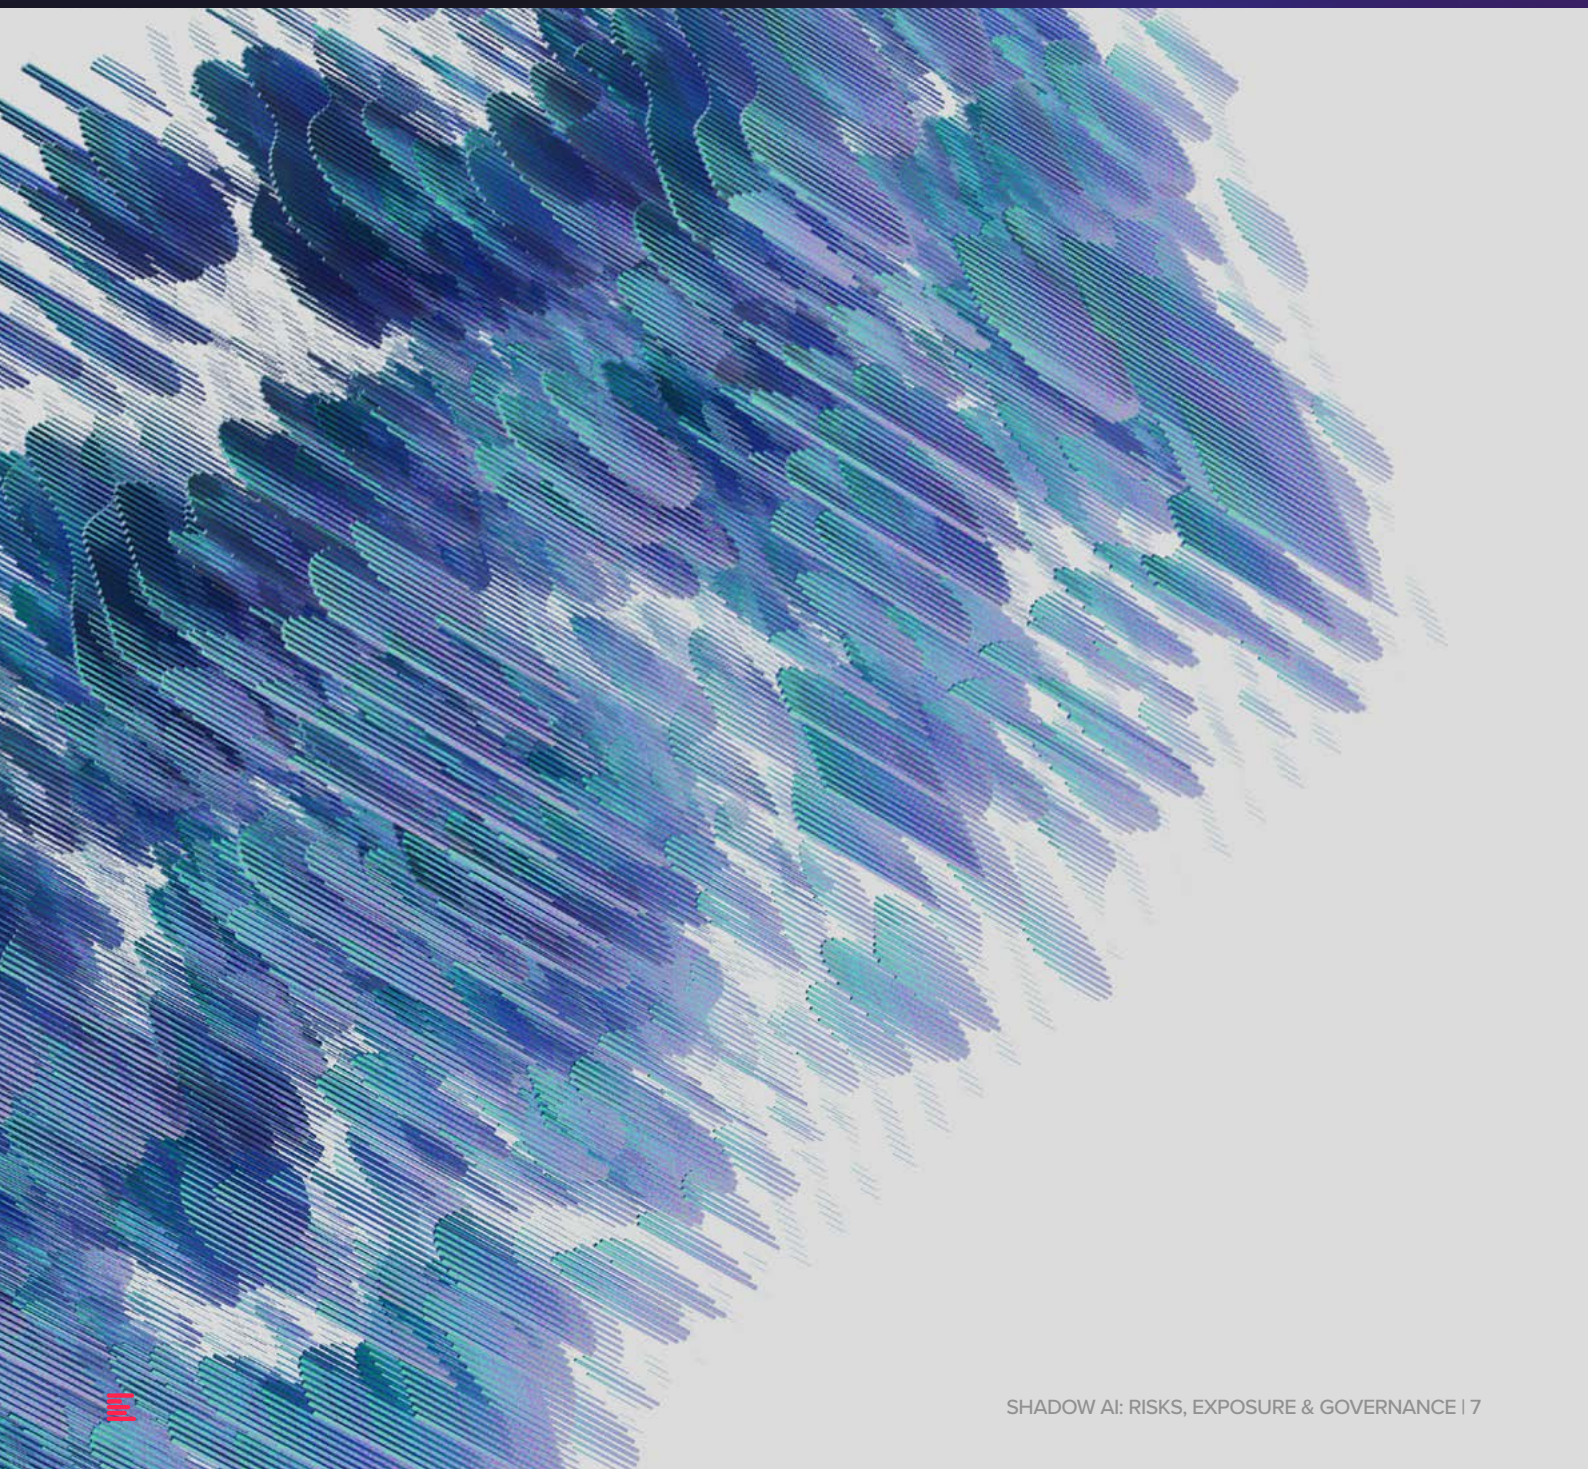

How Shadow AI Manifests Across Tools & Systems

Standalone LLM Tools

General-purpose AI assistants used directly by employees for writing, analysis, coding, and problem-solving. Increasing adoption of desktop AI clients may introduce new risks as these applications interact more directly with local files and system resources.

EXAMPLES:
ChatGPT, Gemini, Claude

Browser Extensions & Overlays

AI-enabled browser extensions that provide assistance embedded directly within the browser. These tools may have broad visibility into user activity and even enterprise SaaS environments. Some tools route requests to multiple AI models, reducing transparency of where enterprise information is ultimately processed.

EXAMPLES:
Grammarly, Sider, Monica

Meeting & Communication Assistants

Meeting assistants capture conversations, creating records of potentially sensitive discussions. Governance challenges include limited control over how meeting data is used and retained, and uncertainty around vendor access to conversations. Even individual licences support integrations with enterprise applications, extending the exposure of meeting data.

EXAMPLES:
Otter.ai, Fireflies.ai, Krisp AI

AI Embedded in SaaS

AI features embedded within approved enterprise software that may be enabled faster than governance teams can fully assess, configure, or monitor. Native governance controls often exist, but visibility and policy enforcement may lag adoption across business units.

EXAMPLES:
Microsoft 365 Copilot, Salesforce Einstein, Notion AI

AI Coding Assistants

AI-powered development tools embedded within IDEs and engineering processes that can generate, review, and modify code while accessing internal repositories, scripts, and technical documentation. Additionally, these tools lower the barrier to software creation, increasing the risk of insecure or insufficiently reviewed code entering enterprise environments.

EXAMPLES:
Cursor, GitHub Copilot, Claude Code

API-Based & Programmatic Use

LLM APIs embedded in internal scripts, applications, or data pipelines enabling scalable and repeatable data processing. As AI use becomes programmatic within operational systems, external data processing may become less visible to governance and security teams than direct employee interaction with standalone AI tools.

EXAMPLES:
OpenAI API, Anthropic API, Google Gemini API, OpenRouter, Ollama

Automation & Agent Platforms

AI-driven automation tools and agent frameworks capable of triggering actions across enterprise systems such as email, CRM, and databases. These systems may connect to enterprise applications through APIs and OAuth permissions, enabling persistent access and semi-autonomous workflows that operate with reduced human oversight.

EXAMPLES:
Zapier Agents, Glean Agents, LangChain

Managing Shadow AI

Organisations need a structured approach to proactively managing Shadow AI. Blanket bans are difficult to enforce and fail to address the underlying drivers of employee behaviour.

This pattern is not new; similar dynamics were seen in earlier waves of enterprise technology adoption, including Shadow SaaS, BYOD, and remote work, where restrictive policies gradually gave way to controlled enablement and governance.

A more effective approach to Shadow AI follows the same trajectory. Instead of trying to remove use, organisations are focusing on providing authorised alternatives that meet employee productivity needs, while embedding controls that reduce data exposure, improve visibility, and enforce policy boundaries.

This requires a layered operating model that combines controlled access to AI, system-level governance and enforcement, and ongoing user awareness to reduce both intentional and inadvertent risk.

Enterprise AI Platforms

Enterprise AI platforms provide a controlled environment for employees to use AI safely, reducing reliance on external tools while maintaining productivity. They act as an approved alternative to Shadow AI by combining usability with baseline governance and data protection.

Recommended Features:

- Data protection controls, including zero data retention and no training on customer inputs
- Identity and access management, with SSO and role-based permissions
- Data isolation and encryption to protect enterprise information
- Logging and auditability of prompts, outputs, and user activity
- Policy controls to restrict sensitive inputs and manage use
- Secure integration with internal data sources under defined access controls

AI Governance & Control Platforms

AI security and control platforms provide visibility, policy enforcement, and risk management across all AI use, including Shadow AI. Unlike traditional security tools that focus on structured data, access, and known patterns, these platforms must govern unstructured interactions, interpret context, and manage dynamic AI behaviour.

Recommended Features:

- Discovery and visibility of AI use across browsers, SaaS, APIs, and endpoints
- Policy enforcement to allow, block, or restrict AI tools and use cases
- Data protection controls to prevent sensitive information from being exposed
- Monitoring and logging of AI interactions for audit and compliance
- Threat protection against risks such as prompt injection, data exfiltration, and attempts to bypass controls
- Control over agent and API activity, including permissions and behaviour enforcement

User Education

Education and user awareness play a supporting role in managing Shadow AI, complementing technical controls and governance frameworks. Many risks arise from well-intentioned behaviour, where employees use AI tools without a clear understanding of how data is processed, stored, or reused. In other cases, employees may deliberately circumvent restrictions by using personal devices.

Providing baseline guidance on acceptable use, data handling, and the risks of AI helps reduce inadvertent exposure and misuse. This ensures that employees can use AI tools more safely, while reinforcing the effectiveness of broader control mechanisms.

BRINGING SHADOW AI INTO THE OPERATING MODEL

This framework can be used as a benchmark to assess where organisations currently stand in their Shadow AI management approach and to identify the capabilities required to strengthen governance, control, and enablement.

FIGURE 3

Managing Shadow AI: Maturity Framework

01.	02.	03.	04.	05.	06.
UNAWARE / UNMANAGED	REACTIVE CONTROL	FRAGMENTED GOVERNANCE	CONTROLLED ENABLEMENT	INTEGRATED GOVERNANCE	OPTIMISED & STRATEGIC
No formal recognition of Shadow AI	Issue identified through incidents or leadership concern	Approved tool lists and basic policies introduced	Enterprise-approved AI environments deployed	AI usage embedded into IT, security, and risk frameworks	AI governance aligned with business strategy
Widespread, invisible usage across the organisation	Initial responses such as bans or access restrictions	Inconsistent enforcement across teams	Guardrails introduced (data controls, access policies)	Monitoring, auditability, & compliance processes established	Controls and tooling embedded into standard operations
	Reliance on existing controls (e.g. network blocking, basic DLP)	Existing security tools partially extended to AI use, with limited effectiveness	Early adoption of AI-specific controls alongside existing security tooling	Combination of existing and AI-specific controls providing consistent oversight	Shadow AI reduced as enterprise alternatives meet user needs
Invisible, uncontrolled usage	**Bans and basic controls**	**Policies without consistency**	**Safe, approved environments**	**Embedded governance and monitoring**	**Strategic, business-aligned use**

CONTROL & TOOLING EVOLUTION

Primarily existing controls
(network, DLP, endpoint)

Mixed approach
Existing tools stretched, early AI-specific tools

AI-specific + integrated controls
Purpose-built AI controls integrated with core security stack

RISK EXPOSURE

High — *Low*

Impact of Shadow AI

Effective management requires action on both sides of the ecosystem – within enterprises to strengthen visibility, controls, and adoption practices, and within vendor ecosystems to ensure AI capabilities are transparent, governable, and compatible with enterprise risk requirements.

IMPLICATIONS FOR ORGANISATIONS

Build a clear view of where AI is already being used

Most organisations underestimate the extent of AI use until they actively look for it. Beyond sanctioned tools, employees use browser-based assistants, extensions, and AI features embedded within SaaS platforms; often without formal approval or visibility. This creates a fragmented usage landscape that is difficult to track through conventional IT monitoring. Establishing a baseline of tools, entry points, and the types of data being exposed is a prerequisite for meaningful governance.

Reassess whether current controls are fit for AI interactions

Traditional security and governance tools were designed for structured systems and predictable data flows. They are less effective in environments where risk emerges through natural language prompts, code generation, and unstructured content exchange. As a result, many organisations can detect activity but struggle to interpret intent or identify sensitive data embedded in AI interactions. The gap between visibility and meaningful control is where most exposure currently sits.

Treat education and governance as part of adoption, not a separate layer

Attempts to manage Shadow AI through standalone training or policy documents tend to have limited impact. Behaviour is more effectively shaped when guidance is embedded directly into how AI tools are introduced and rolled out across the organisation. Similar changes in security awareness show that embedding guidance into day-to-day workflows proved more effective than standalone, periodic training. The objective is not just awareness, but consistent decision-making in everyday use.

Extend oversight to AI-driven integrations and automation

Shadow AI risk is shifting from individual tool use to embedded automation. APIs, scripts, and agent-based workflows can maintain persistent access to enterprise systems through credentials and permissions, often operating outside direct human visibility. As AI becomes more machine-to-machine, governance needs to extend beyond user activity to include integration approval, access management, and continuous monitoring of automated processes.

IMPLICATIONS FOR TECHNOLOGY PROVIDERS

Be explicit about how AI is used and where data goes

As AI becomes embedded across products, enterprises expect clarity on what data is processed, where it is sent, and how it is stored or reused. This includes AI features inside existing tools, not just standalone models. Lack of transparency around data flows and model routing is becoming a friction point in enterprise adoption.

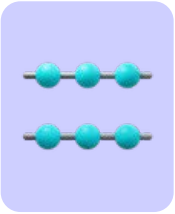

Build governance controls into the product, not around it

Enterprise customers are evaluating AI capabilities based on whether they can be governed. This includes controls for data retention, training behaviour, access permissions, and audit logs. Providers that treat governance as an external requirement will face slower enterprise uptake.

Give customers control over AI feature activation

AI functionality is being introduced inside established SaaS and developer tools by default or through rapid release cycles. Enterprises need the ability to selectively enable, restrict, or disable AI features by user group, region, or use case.

Improve visibility across models, APIs, and integrations

As AI systems rely on multiple models and third-party services, enterprises need clearer visibility into how requests are routed and processed. This includes dependencies across APIs, intermediate processing layers, and external model providers.

Support integration with enterprise monitoring and policy tools

AI tools are now expected to operate within existing security and governance ecosystems. This requires compatibility with enterprise logging, monitoring, and policy enforcement systems, as well as the ability to export usage data for audit and compliance purposes. Closed or unobservable systems are difficult to adopt.

About the Authors

Darian Bird

Principal Advisor, Ecosystm

Darian helps businesses navigate transformation, providing insight into cloud, automation, data management, infrastructure, and telecommunications. He has spent two decades advising leaders on leveraging technology to enter new markets, improve client experiences, and enhance service delivery, with a particular focus on infrastructure and strategic resilience. At Ecosystm, Darian conducts in-depth research on cloud, data, and AI sovereignty regulations across the Asia Pacific, exploring their implications for enterprises and technology providers.

Sash Mukherjee

VP Industry Insights, Ecosystm

Sash is a research industry veteran with nearly 25 years of experience in analysis, writing, and training across sectors including Public Sector, Healthcare, Education, and Insurance. She plays a key role in shaping Ecosystm's research strategy, developing research-backed thought leadership that guides enterprises on industry trends and supports tech vendors' targeted go-to-market strategies. She is also involved in delivering consulting projects and custom engagements, and is a regular speaker and panelist at industry events.

ecosystm.

Ecosystm is a leading technology market analyst and advisory firm that helps stakeholders navigate innovation in the digital economy through data, insights, and expertise. We connect enterprises, technology companies, digital-native founders, investors, and policymakers to enable informed decision-making. With ongoing research and access to top analysts and strategic advisors, we empower business planning, go-to-market activities, thought leadership, and innovation strategy consulting. Visit ecosystm.io