SOPHOS

MSP GUIDE

Building an AI governance practice: A guide for MSPs

As organizations expand their adoption and use of AI, they need effective governance to unlock its full value without exposing themselves to unacceptable cyber risk. This guide provides a blueprint for MSPs to develop high-value, repeatable AI governance services that meet growing customer demand, starting by building the same capabilities within their own business.

The AI governance gap

Organizations are adopting artificial intelligence (AI) faster than they can secure it, with Gartner predicting that by 2030 more than 40% of enterprises will experience security or compliance incidents linked to unauthorized shadow AI.[1]

In many cases, AI tools and agents are already deeply woven into business operations, connecting with sensitive data, trusted identities, and critical applications. As adoption accelerates, governance is failing to keep pace, resulting in a growing attack surface that organizations struggle to uncover, control, and protect.

The business challenge has moved beyond deciding whether employees should use AI tools. Organizations need help capturing AI's value while keeping adoption aligned with their risk appetite, policies, and regulatory obligations. Without effective governance, gains in productivity, efficiency, and growth can easily be offset by unmanaged exposure, weakening returns, and longer-term business risk.

Many organizations lack the expertise or capacity to address that exposure alone, looking instead to their service partners for this kind of cybersecurity leadership. With 84% of managed service providers (MSPs) expecting demand for CISO-level services to increase over the next 12 months,[2] AI governance is an organic extension of that existing relationship.

Customers need partners that can help them understand their AI risk profile, establish effective governance, and continuously adapt as AI technologies evolve. For MSPs, this presents an opportunity to:

01 Open or deepen strategic conversations around governance, risk, compliance, and cyber resilience

02 Develop repeatable advisory services with recurring revenue potential

03 Help customers put governance into practice with solutions that discover, understand, and control AI use across their environments

The opportunity is clear. The question is whether MSPs are ready to seize it.

40%

Of enterprises will experience incidents linked to unauthorized shadow AI.

GARTNER

84%

Of MSPs expect demand for CISO-level services to increase in the next 12 months.

SOPHOS MSP PERSPECTIVES

Contents

The AI governance gap 2

Key definitions 4

Understanding the market need for AI governance 5

The MSP AI governance readiness gap 8

Building an AI governance practice 11

Starting your AI governance journey 16

Bring AI out of the shadows with Sophos 17

Key definitions

Understanding these terms will help you navigate this guide and put the information into practice.

AI governance

The policies, roles, processes, and controls used to ensure AI is adopted and operated responsibly, securely, and in line with business and regulatory requirements.

AI agent

An AI system programmed to independently plan and perform tasks, interact with applications, and use data or tools on a user's or organization's behalf.

AI model

A core mathematical system or algorithm trained on data to recognize patterns, generate outputs, make predictions, or support automated decision-making.

AI tool

An application or software service that uses one or more AI models to help users complete specific tasks.

AI integration

A connection that embeds AI capabilities into another application, system, or workflow so it can access organizational data or trigger actions automatically.

Shadow AI

The adoption of AI tools, browser extensions, or features by employees without formal approval, visibility, or oversight from IT and security teams.

Prompt injection

A security vulnerability or attack where malicious or engineered instructions manipulate an AI model into ignoring system rules, bypassing guardrails, or executing unauthorized actions.

Human-in-the-loop (HITL)

A governance approach requiring human review, validation, or intervention at designated points in an AI-driven workflow.

Retrieval-augmented generation (RAG)

An architecture that improves generative AI responses by querying and retrieving context from external authoritative databases before generating an answer.

Model drift

The degradation or shift in an AI model's accuracy, safety, or operational behavior over time due to changing real-world data patterns, user inputs, or environment changes.

Understanding the market need for AI governance

AI adoption is creating value for organizations everywhere, but it is also introducing risk faster than many of them can establish visibility, ownership, and control.

The state of AI adoption today

AI has rapidly moved into everyday business use, bucking the conventional technology adoption curve. Today, the focus of the security community is no longer whether organizations are adopting AI, but how deeply it has entered their operations and whether governance has kept pace.

According to McKinsey's *The state of AI in 2026: On the road to ROI*,[3] nearly 90% of respondents report regular AI use within at least one business function. Adoption is also deepening, with 44% saying AI is scaling across their enterprise, up from 38% a year earlier.

Among smaller businesses, AI use is growing quickly, but deeper integration remains limited. In the U.S., 76% of small businesses surveyed by Goldman Sachs[4] in its *10,000 Small Businesses Voices* are using AI, yet only 14% have fully integrated it into their core operations. In the U.K., NatWest[4] found that 44% of SMEs and midmarket businesses use AI, but only 6% are doing so at a transformational level. Among those using AI, 29% remain in the earliest stages of piloting and awareness building.

Innovation always outpaces security at first, but with AI, the window of risk is far greater. As AI tools, agents, and integrations become embedded in business processes at record speed, often without the governance needed to manage them safely, new, unmanaged paths to data exposure, compromise, and business disruption are being formed. For threat actors, every ungoverned tool, agent, and integration becomes another potential target.

~90%

Of organizations report regular AI use in at least one business function

MCKINSEY

44%

Say AI is scaling across their enterprise, up from 38% a year earlier

MCKINSEY

76%

Of U.S. small businesses are using AI, yet only 14% have fully integrated it

GOLDMAN SACHS

6%

Of U.K. SMEs and midmarket businesses have reached transformational adoption

NATWEST

Organizational AI security & operational risks

With every AI tool, feature, agent, or integration woven into their everyday operations, an organization's attack surface expands far beyond traditional enterprise boundaries. Existing threat vectors across data, identities, applications, endpoints, networks, and cloud environments are now intertwined with novel AI entry points. These include user prompts, generated responses, AI orchestration layers, RAG vector stores, and third-party plugins and skills. The result is a deeply interconnected and sensitive web of exposure that is uniquely difficult to identify, inventory, and control due to its broad and decentralized distribution across the organization. As a result, organizations now face a whole set of new risks to govern.



Interconnected operations

AI agents act as connective tissue across enterprise data, credentials, APIs, and cross-application workflows. While this broad integration drives automation, it introduces systemic risk: a single prompt injection, agentic drift, or compromised API connection can cascade downstream through connected workflows. These dependencies exponentially expand the blast radius of security incidents, allowing a single point of compromise to escalate into enterprise-wide operational failure.



Shadow AI and inventory blind spots

Unsanctioned AI use, from employee-deployed agents to AI capabilities introduced via routine SaaS feature updates, creates critical governance gaps. Because AI adoption often bypasses traditional procurement pathways, IT and security teams easily lose visibility.

Without a comprehensive inventory of active AI models and their permission scopes, data inputs, and execution capabilities, organizations cannot maintain an effective defensive perimeter or establish zero-trust boundaries.



Data exfiltration

Sensitive corporate data can leave secure systems through unsecure AI prompts, file attachments, and automated API fetches. As autonomous agents retrieve and pass contextual data between disparate tools and third-party AI models, data boundary enforcement degrades, significantly raising the risk of data leakage, unauthorized retention, and regulatory non-compliance.



Privilege escalation and autonomous action

AI agents can inherit user credentials, API keys, and elevated execution permissions to perform complex tasks. When paired with autonomous decision-making capabilities, these excessive permissions allow agents to read, modify, or delete sensitive data without HITL oversight. Prompt injection and exploitation or logic errors can trigger unauthorized execution of unsecure actions too fast for security teams to intervene.



Altered cyberattack economics

AI has increased the speed, scale, and scope of cyberattacks. Threat actors can now use AI to accelerate reconnaissance and exploitation, automate activity across multiple targets, and tailor campaigns to specific organizations and users. When attacks reach AI-connected systems, the effects can also spread across the data, identities, applications, and workflows those systems have access to, increasing the reach, velocity, and potential impact of a compromise.

Why customers need help

Organizations that cannot systematically evaluate how risk escalates based on *who or what* is interacting with AI, which sensitive systems or datasets are exposed, and what autonomous actions are executed, will struggle to keep that risk within acceptable limits. Without the right controls, their exposure will continue to grow.

The central challenge is that rapid AI adoption has outpaced many organizations' ability to strategically identify and enforce those controls. As they struggle to keep up, they increasingly need specialist, external support. This creates an inroad for MSPs and other channel partners. Since these specialists already manage the core IT, security, network, and compliance infrastructure that AI runs on, they are naturally suited to address customer AI governance gaps by:



Leveraging existing system access:

MSPs already have visibility into the endpoints, cloud environments, networks, and identity frameworks where customers' AI systems and workflows operate.



Bridging technical and executive gaps:

MSPs excel at translating complex technical risks into practical, business-level policies for non-technical stakeholders.



Providing continuous, scalable oversight:

Rather than delivering a one-time audit, MSPs can embed AI discovery, guardrail monitoring, and risk assessment directly into their ongoing managed security services.



Operationalizing industry frameworks:

Customers struggle to interpret abstract guidance (like NIST or ISO standards); MSPs can turn guidelines into turnkey, repeatable service offerings.

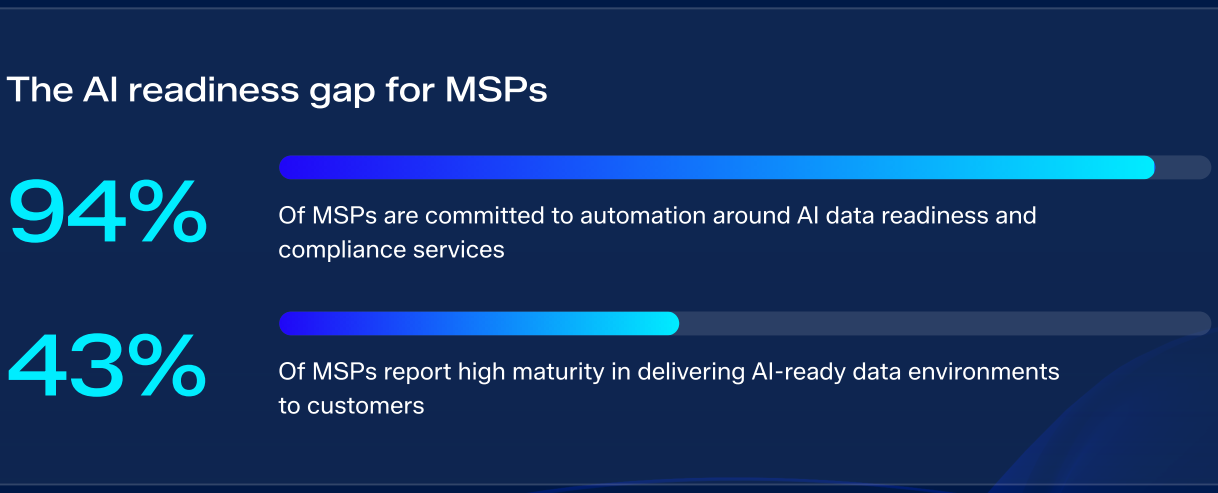
By stepping in to bridge the gap between rapid AI adoption and effective AI governance, MSPs can address their customers' critical weaknesses while expanding their own portfolio into high-margin advisory and managed security practice areas.

The MSP AI governance readiness gap

MSPs should increasingly recognize AI as a strategic capability and a potential source of new customer revenue. However, recognizing the opportunity and being ready to deliver against it don't always go hand-in-hand. A repeatable service needs a common method for discovery, recommendations, implementation, and follow-up.

Research conducted by Omdia[5] found that 94% of surveyed MSPs are committed to automation around AI data readiness and compliance services, but only 43% report high maturity in delivering AI-ready data environments to customers. This means that more than half therefore lack the maturity to deliver these services at scale.

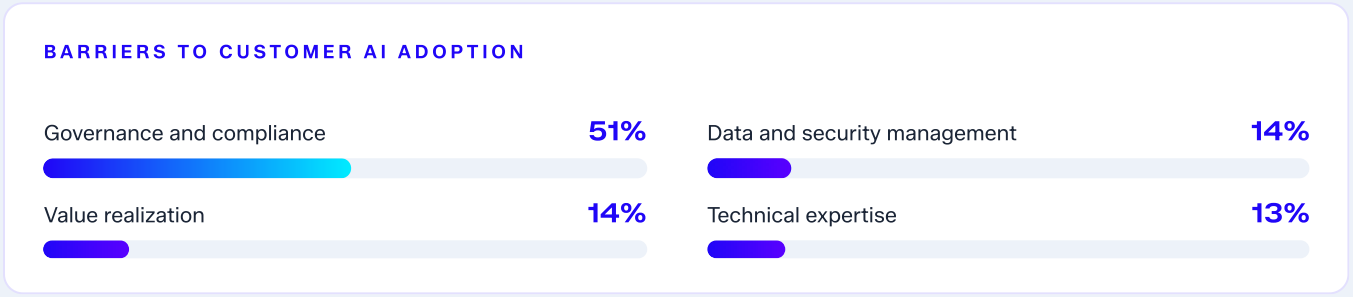
The AI readiness gap for MSPs



The tech talk trap

Governance and compliance are the most commonly cited barriers to customer AI adoption. In the same Omdia research,[5] 51% of MSPs identified these as their primary obstacle, compared with 14% citing data and security management, 14% value realization, and 13% technical expertise.

This creates an opening for MSPs, but only if they engage customers at the right level. Conversations that focus narrowly on which tools to approve or block rather than how AI risk should be governed across the entire organization reduce the opportunity to a product-by-product transaction.



AI governance in practice requires a deeper analysis of customers' business objectives, cybersecurity strategy, risk posture, privacy needs, and compliance demands. MSPs must be able to explain how risk changes as different AI models are introduced, sensitive data moves through them, use cases expand, access broadens, and regulatory obligations evolve, as well as how those factors change depending on the environment. To help customers overcome the AI governance barrier, MSPs must ultimately move the dial from "Which AI tools should we allow?" to "How do we identify and control AI use in context?"

The diagram below illustrates the journey from MSPs simply operating AI technology for customers to becoming their strategic AI security partner. Effective AI governance can be the vehicle that moves you along this curve, taking the conversation beyond tools and controls into business objectives, security strategy, and cyber resilience goals. By helping customers transform their approach to AI from the ground up, MSPs can build their role around continuous value and achieve recurring revenue.



Guidance may be difficult to scale

From tool selection to compliance, organizations without the AI knowledge capital internally will need their security partners to be the subject matter expert in AI security and governance. The issue is that, although half of customers already look to their MSP for this form of leadership,[6] few MSPs have transitioned from ad-hoc advice to value-creating CISO-level consultancy. So long as the guidance MSPs provide remains informal, reactive, and dependent on individual expertise, they will struggle to keep up with demand.

Any MSP that wants to be at the forefront of AI governance must adopt a model that enables them to deliver CISO-level outcomes at scale, demonstrate measurable value, and create pipeline from every interaction.

MSPs must overcome the same governance gaps

MSPs are putting AI to work within their own businesses. Omdia's *MSP Analysis January 2026*[7] identifies that MSPs are increasingly adopting AI internally as a strategic capability and identifies that:

> The most successful MSPs deploy AI first within their own operations to drive efficiency, improve service quality, and build repeatable intellectual property. This internal adoption then becomes the foundation for monetized customer-facing AI services.
>
> **OMDIA, MSP ANALYSIS JANUARY 2026**

Like their customers, MSPs may still be discovering which tools and agents are in use, what data they access, who owns the associated risk, and whether everyday practices align with policy. Working through those questions internally gives MSPs direct experience of the practical challenges involved and enables them to turn that into a repeatable blueprint. And by revealing and testing that approach within their own operations first, they can build the knowledge, processes, and intellectual property needed to guide customers through the same journey, thereby developing a credible, scalable AI governance service.

Building an AI governance practice

Effective AI governance translates complex technical and regulatory requirements into a structured operational framework. Through continuous discovery, risk mapping, control alignment, and accountability, customers can turn abstract requirements into practical controls and maintain governance as their AI use evolves.

Gartner describes this systematic approach as AI Trust, Risk, and Security Management[8] (AI TRiSM). For MSPs, AI TRiSM provides a useful way to frame the capabilities customers increasingly need: visibility into AI use, risk assessment in business context, proportionate controls, clear accountability, and continuous monitoring as technologies and regulatory requirements evolve.

The foundations of governed AI

MSPs must understand the principles that underpin effective AI governance. The questions below are designed to help you identify where to strengthen your own knowledge before attempting to develop your AI governance practice.

VISIBILITY	☐ What AI models, tools, agents, and integrations are in use?	☐ Which data, identities, systems, and workflows can they access?
CONTEXT	☐ What business purpose does each AI use case support?	☐ How could misuse, compromise, or failure affect our organization?
CONTROL	☐ Which AI uses should be approved, restricted, or prohibited?	☐ What guardrails and human oversight are proportionate to the risk posed by each AI use case?
OWNERSHIP	☐ Who owns each AI use case and its associated risk?	☐ Who approves exceptions and responds when controls fail?
IMPROVEMENT	☐ How will our organization track changes and demonstrate success?	☐ What changes should trigger reassessment?

MSP AI governance offering blueprint

Effective AI governance is a continuous cycle: discover AI use, assess risk in context, map data and system dependencies, apply proportionate controls, and monitor change over time. MSPs must first establish this cycle within their own operations, giving them the practical experience to turn it into a repeatable customer service.

This section provides a blueprint for putting that cycle into practice, strengthening your own AI governance while building the foundations for consistent, repeatable services for your customers.

OVERVIEW

0 **Ownership**	Assign key roles to the relevant stakeholders needed to develop an effective practice.
1 **Discovery & visibility**	Assess internal readiness gaps as well as customer demand, AI adoption, and governance needs.
2 **Strategy & proposition**	Define your strategy, portfolio alignment, and customer value proposition.
3 **Design & operationalization**	Package capabilities into standardized, repeatable services with clear delivery models.
4 **Internal deployment ("dogfooding")**	Test services internally to validate controls, efficiency, and delivery expertise.
5 **Delivery & monetization**	Launch scalable customer services with continuous assessment, improvement, and reporting.

MSP AI governance offering blueprint

PHASE

0

Ownership

Instruct your AI governance team

AI governance sits at the intersection of security, legal, operations, business strategy, and customer engagement. Bringing together the right stakeholders from the outset creates the alignment and expertise needed to develop an effective practice.

KEY STAKEHOLDERS

- **Executive sponsor:** Leads on strategic objectives and delivery.
- **CISO/security lead:** Provides risk management, compliance, and policy oversight.
- **Service delivery lead:** Ensures recommended processes, technologies, and controls can be operationalized and scaled.
- **Sales/account leadership:** Shapes commercially viable services around the strategy.

PHASE

1

Discovery & visibility

Establish your current state

Your AI governance practice must have a strong, evidence-based foundation. This begins by assessing where and how AI is currently being deployed and used, both within your own business and across your customer base.

KEY QUESTIONS

- **Internal usage & exposure:** What AI tools, agents, and embedded application features are currently in use, and what data, identities, or system permissions do they interface with?
- **Governance gaps:** Where does AI adoption and use not align with existing security policies, controls, and procurement workflows?
- **Customer readiness:** What AI capabilities are your customers actively adopting, what risks are they concerned about, and what internal visibility gaps are they already aware of?

STRATEGIC DELIVERABLE	**AI baseline assessment:** A unified, high-level view of your own internal operational readiness and customer demand highlighting the immediate risks and opportunities to be addressed.

PHASE

2

Strategy & proposition

Define your position and commercial value

Develop your established baseline and initial learnings into a strategy that reflects your business' risk appetite and define the proposition your business needs and that you want to take to market.

KEY QUESTIONS

- **Risk tolerance:** What levels of AI autonomy, data sharing, and system integration are acceptable within your own business and for your target market?
- **Service alignment:** How does AI governance fit into your existing managed security, compliance, and CISO-style advisory portfolio?
- **Value proposition:** What specific business outcomes (e.g., risk reduction, regulatory compliance, secure growth) must your AI governance services deliver to for it to be relevant and competitive?

STRATEGIC DELIVERABLE	**AI strategy and service proposal:** A documented internal AI strategy alongside a structured, commercial value proposition for customer-facing governance services.

PHASE

3

Design & operationalization

Productize your governance expertise

Turn your strategic framework and value proposition into standardized modules, testing them within your own organization before developing them into repeatable customer services.

KEY QUESTIONS

- **Service tiering:** How will you structure your offerings internally and externally (e.g., initial AI risk assessments, ongoing guardrail monitoring, governance consultancy)?
- **Framework selection:** Which recognized industry standards (such as NIST AI RMF or ISO/IEC 42001) must your delivery methodology account for?
- **Operational readiness:** What skill sets, tooling, and reporting templates are needed to deliver these services at scale?

STRATEGIC DELIVERABLE	**Service catalog:** A packaged, repeatable portfolio of AI governance capabilities complete with defined internal and external use cases.

PHASE

4

Internal deployment ("dogfooding")

Validate modules within your own business

Test your AI governance service modules internally. Capture data on the effectiveness of your delivery, implement changes, and build authentic reference cases.

KEY QUESTIONS

- **Business impact:** How do your AI governance services affect day-to-day employee productivity and workflow efficiency?
- **Control verification:** Are your controls and guardrails successfully catching unapproved AI usage or high-risk data exfiltration, and where could there be gaps?
- **Intellectual property:** What unique insights, playbooks, and best practices can be captured from your internal adoption to enhance customer delivery?

STRATEGIC DELIVERABLE	**Internal governance practice:** A scoped internal working team and environment that serve as a living testbed to demonstrate your capability to prospective customers.

PHASE

5

Delivery & monetization

Scale customer engagements and maintain continuous oversight

When your internal practice is fully operational, you can begin to roll out your structured AI governance offerings to the market, helping customers continuously manage AI security while generating high-margin advisory and managed service revenue.

KEY QUESTIONS

- **Go-to-market execution:** How will you integrate AI risk assessments into your standard customer account reviews and CISO-style strategic discussions?
- **Continuous monitoring:** How will you track evolving customer AI usage, emerging regulatory mandates, and model drift over time?
- **Client value demonstration:** What metrics and executive reporting formats will prove ongoing risk reduction and customer ROI?

STRATEGIC DELIVERABLE	**Managed AI governance practice:** A functional, market-ready AI governance practice that can deliver a recurring revenue stream across your customer base.

Starting your AI governance journey

Rapid AI adoption across data, identities, applications, and autonomous workflows has created a critical governance gap for organizations. While many are eager to harness AI, their ability to manage and secure it has fallen behind.

Because MSPs already manage the underlying IT and security parameters where AI operates, they are uniquely positioned to solve this challenge. By implementing an AI governance practice within their own business first, MSPs can close internal readiness gaps, build intellectual property, and turn customer complexity into a high-margin, managed advisory practice.

What to do next

1 Assess your own AI governance and readiness

Your customers' AI governance starts with your readiness to deliver it. To meet the market need, MSPs must establish an evidence-based baseline within their own businesses. Use your operations as a testing ground to audit active AI tools, test technical guardrails, and refine policies, turning your internal risk management into your strongest commercial proof point.

2 Prepare to talk to customers

Initiate strategic conversations with customer leadership to uncover hidden visibility gaps and quantify their exposure. Open these discussions with three high-level discovery questions:

Visibility:
What AI tools, embedded application features, or autonomous agents are currently active across your workforce?

Exposure:
What sensitive data, credentials, and core systems are those AI technologies allowed to access and interact with?

Governance:
Where do your existing security controls and internal policies fall short of actual employee AI usage?

3 Ensure you have the right solutions to deliver AI governance

Evaluate the platform capabilities, technical visibility tools, and strategic frameworks needed to back up your advisory services. Partnering with purpose-built security platforms enables you to automate AI discovery, enforce technical controls, and deliver continuous managed governance at scale.

Bring AI out of the shadows with Sophos

SOPHOS AI DEFENSE

Sophos AI Defense helps you understand and control AI use across your customers' environment, combining visibility, risk insight, and clear enforcement in one place. Available to eligible Sophos EDR, Sophos XDR, and Sophos MDR customers through Sophos Fusion, it extends their existing investment into a coordinated approach to AI security.

- **Discover sanctioned and shadow AI use**, revealing activity that may otherwise remain outside the customer's visibility.
- **Assess risk and protect sensitive data** based on how AI is being used across the environment.
- **Turn governance policies into enforceable guardrails** without unnecessarily restricting approved AI use.
- **Bring AI-related detection and response into existing Sophos workflows**, without a separate, heavyweight platform.

Explore Sophos AI Defense →

SOPHOS CISO ADVANTAGE

Sophos CISO Advantage gives MSPs an inroad to AI governance, helping you turn the CISO role you are already delivering into a scalable service. Shape and manage your customers' cyber programs at scale with AI-accelerated assessment, analysis, and reporting. Demonstrate value, strengthen relationships, and create new revenue opportunities.

- **Assess and contextualize** AI risk within the customer's wider security, compliance, and business environment.
- **Translate posture and compliance gaps** into clear priorities, business-ready reporting, and improvement roadmaps.
- **Build the case for further action**, including the operational visibility and controls provided by Sophos AI Defense.
- **Monetize governance** through repeatable assessment, reporting, strategy, and improvement services that generate recurring revenue.

Explore Sophos CISO Advantage →

Together, Sophos AI Defense and Sophos CISO Advantage give MSPs a route from identifying, prioritizing, and communicating AI risk to implementing the visibility and controls needed to reduce it.

SOPHOS

Want to become a Sophos Partner? Register here.

United Kingdom and Worldwide Sales

Tel: +44 (0)8447 671131

Email: sales@sophos.com

North America Sales

Toll Free: 1-866-866-2802

Email: nasales@sophos.com

Australia and New Zealand Sales

Tel: +61 2 9409 9100

Email: sales@sophos.com.au

Asia Sales

Tel: +65 62244168

Email: salesasia@sophos.com

1 - https://www.gartner.com/en/newsroom/press-releases/2025-11-19-gartner-identifies-critical-genai-blind-spots-that-cios-must-urgently-address0
2 - http://sophos.com/en-us/resources/report/2026-msp-perspectives-report
3 - https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai#/
4 - https://www.goldmansachs.com/pressroom/press-releases/2026/small-businesses-embrace-ai-but-need-training-and-support-to-fully-harness-it
5 - https://www.globenewswire.com/news-release/2026/04/09/3271007/0/en/research-from-avepoint-and-omdia-reveals-governance-and-compliance-as-the-leading-ai-adoption-barrier-among-msps.html
6 - https://www.sophos.com/en-us/resources/report/2026-msp-perspectives-report
7 - https://omdia.tech.informa.com/om142050/msp-trends-and-predictions-2026
8 - https://www.gartner.com/en/articles/ai-governance-trism

GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved.

© Copyright 2026. Sophos Ltd. All rights reserved. Registered in England and Wales No. 2096520, The Pentagon, Abingdon Science Park, Abingdon, OX14 3YP, UK. Sophos is the registered trademark of Sophos Ltd. All other product and company names mentioned are trademarks or registered trademarks of their respective owners.

2026-11-18 (MC)