DARKTRACE

Darktrace / SECURE AI™

Secure your AI adoption with Adaptive AI that delivers real-time visibility and control across prompts, agents, development environments, and Shadow AI.

88%

of organizations
are using AI in at least one business function[1]

80%

of Darktrace customers
now use generative AI services[2]

44%

of security professionals
are very or extremely concerned about AI agents[3]

AI adoption is outpacing enterprise control

AI has entered the enterprise at a rapid pace that is only accelerating. While adoption and innovation move forward, security is still catching up. Organizations now face an AI security problem, with risk compounding quietly across workflows, tools, data, and decisions, often without any visibility at all.

The adaptive and unpredictable nature of AI means effective AI security requires a deep understanding of what "normal" looks like across humans, AI systems, and agents, so that the earliest signs of drift, misuse, or compromise can be caught before they escalate.

BUSINESS BENEFITS

Unified AI visibility
See AI interactions across enterprise GenAI, agents, development activity, and Shadow AI through one view of behavioral intelligence.

Real-time risk detection
Detect subtle misuse, drift, and emerging risk in prompts and agent behavior as it happens, without relying on predefined rules.

Policy-aligned governance
Map your own organizational policy to AI activity, so governance is aligned to business priorities and evaluated against how AI is actually used.

Confident AI adoption at scale
Expand AI use across the business with the visibility, understanding, and oversight needed to manage risk at scale.

1: McKinsey State of AI 2025
2: Darktrace Research
3: Darktrace State of AI Cybersecurity 2026

Bringing behavioral security to your AI adoption

Powered by Adaptive AI, Darktrace / SECURE AI applies more than a decade of behavioral understanding to secure how prompts, agents, and AI systems are used across your business

Rather than relying on predefined rules or signatures, it learns your organization's operational patterns and surfaces the misuse, misconfiguration, and drift that traditional, rule-based controls are not designed to detect.

Delivered through the Darktrace Platform Portal, Darktrace / SECURE AI brings AI interactions into a single view, helping teams understand intent, assess risk, protect sensitive data, and enforce policy across both human and AI agent activity.

BUILT FOR AI IN YOUR ENTERPRISE

Enterprise GenAI	AI Agents
Low-Code	High-Code / Cloud AI
Shadow AI	SaaS
MCP	SASE

How Darktrace secures your AI

Darktrace's Adaptive AI already interprets ambiguity, uncovers subtle intent, and spots drift before activities or events become dangerous across your enterprise.

Extending that same behavioral approach to AI is a natural next step, not a leap, building on the Darktrace Behavioral Defense Platform™ that already protects people and infrastructure.

Discover and control Shadow AI

Darktrace helps you surface unsanctioned or unexpected AI activity wherever it appears, including unauthorized low-code agent building and hidden agentic infrastructure such as locally deployed MCP servers, by combining native / SECURE AI visibility with network, cloud, endpoint, and SASE data sources. Policy and response paths appropriate to the environment then help contain data exposure and guide users toward sanctioned options.

Monitor the prompts driving GenAI agents and assistants

Darktrace / SECURE AI provides real-time inspection of prompts, sessions, and responses across enterprise AI like Microsoft Copilot and ChatGPT Enterprise, low-code environments like Copilot Studio, high-code environments like Amazon Bedrock, and SaaS suppliers like Salesforce and Microsoft 365. Behavioral analytics can surface subtle deviations in AI activity that may signal emerging risk, helping organizations detect potential issues such as sensitive data exposure, jailbreak attempts, and other unsafe or unexpected behavior.

Secure AI agent identities in real time

You can discover agent identities across suppliers and environments, map their permissions, roles, and connections, and audit behavior across SaaS, cloud, network, endpoint, OT, and email. This real-time audit view gives organizations a clear understanding of their AI agents, including where they are deployed, how they are being used, and what permissions they have, helping teams identify potential risk and maintain oversight at scale.

Evaluate AI risks in development and deployment

With Darktrace / SECURE AI you gain visibility into the identities, permissions, and access being assembled across low-code builders like Copilot Studio, high-code platforms such as Amazon Bedrock, and internal labs, with posture management that flags misconfigurations and anomalous build events before an agent reaches production. / SECURE AI then connects that build-time insight to runtime reality by linking development context to prompt oversight.

Choose the protection you need

Darktrace / SECURE AI is available in packages designed to support different levels of AI security maturity, following the same Core, Advanced, and Complete adoption path as the rest of the Darktrace Behavioral Defense Platform.

CORE

Includes AI prompt analysis and integrated Shadow AI management to provide visibility into sanctioned AI usage and associated risk, as well as policy alignment mapping.

ADVANCED

Extends Core with deep Shadow AI management, AI Agent identities and actions, and AI agent development risk management for deeper oversight of AI activity, agent behavior, and risk.

"Immediately, we became privy to the information that we'd been lacking all along: **who is using what** within our organization and **what are they doing with the data.**"

■ VP of IT, US Manufacturing

■ **About Darktrace**

Darktrace is a global leader in AI for cybersecurity that keeps organizations ahead of the changing threat landscape every day. Founded in 2013 in Cambridge, UK, Darktrace provides the essential cybersecurity platform to protect organizations from unknown threats using AI that learns from each business in real-time. Darktrace's platform and services are supported by 2,300 employees who protect nearly 10,000 customers globally. To learn more, visit darktrace.com.

North America: +1 (415) 229 9100 | Europe: +44 (0) 1223 394 100 | Asia-Pacific: +65 6804 5010 | Latin America: +55 11 4949 7696

darktrace.com | info@darktrace.com

© 2026 Darktrace Holdings Limited. All rights reserved.