Mage Data™ Launches Data Security and Privacy for AI, Extending Enterprise Data Protection Across the AI Lifecycle

Mage Data extends enterprise data protection across the AI lifecycle with guardrails for training data, usage, masking, development, and monitoring.

NEW YORK, NY, UNITED STATES, August 31, 2026 /EINPresswire.com/ -- Mage Data, a Leader in Data Security Platforms (KuppingerCole), Champion in Test Data Management (Bloor Research), and a notable vendor in Forrester's The Data Security Platforms Landscape, Q3 2026, today announced Data Security and Privacy for AI, extending its data protection platform into AI environments. The release applies consistent policy to data entering an AI system, data leaving one, and the activity in between — across training pipelines, public generative-AI assistants, custom agents, and embedded copilots.

Why Existing Controls Fall Short in AI Environments
Enterprise data protection was built for queries, tables, and applications. AI systems break those assumptions all at once: they answer in natural language rather than rows, they can be questioned indirectly and iteratively, they are often bought rather than built, and they copy sensitive data long before anything reaches production. As a result, conventional controls struggle to:

- Protect data across the extracts, notebooks, feature stores, and evaluation sets that multiply during model development
- Prevent employees pasting customer information, contract terms, or source code into public AI tools
- Tie what an agent can retrieve to what its user is entitled to see — including in SaaS assistants whose UI, agent logic, and model cannot be modified
- Catch inference attacks, where a harmless-looking answer becomes sensitive in light of the question behind it

Five Guardrails Across the AI Lifecycle
Data Security and Privacy for AI builds on Mage Data's foundation of Sensitive Data Discovery and Classification, Static Data Masking (SDM), Dynamic Data Masking (DDM), and Activity Monitoring, extending each across AI workloads — wherever they run. The capabilities delivered include:

• Training Data Guardrails discover PII, PHI, and NPI across structured and unstructured data and protect it at the right stage of your AI pipeline — masking at the source, securing data as it's provisioned into pipelines, or embedding controls in your code via SDKs — so data is governed and protected before any model sees it.

• AI Usage Guardrails inspect all employee prompts and file uploads to public gen-AI tools and mask sensitive content before it leaves the device — masking, never blocking, so protection doesn't push usage underground.

• Dynamic Data Masking for AI masks, redacts, generalizes, or blocks a response based on who asked, what was asked, and what the AI is about to say — protecting purchased copilots or agentic apps without altering them.

• AI Development Guardrails let teams build least privilege into their own agents through Mage SDKs and an MCP Server — each session's tools and data access are scoped to the end user's entitlements, so an agent can't hand a user data they were never authorized to see.

• Activity Monitoring for AI records every AI interaction — who asked what, which tool, use of unsanctioned AI tools, what data was masked, any overrides, and the policy outcome — feeding dashboards, end-user reports, and real-time alerts.

Enterprises can extend policies already deployed on the Mage Data platform directly to their AI workloads — no parallel policy framework, no separate compliance regime for AI.

Leadership Perspectives on the Launch

"The conversation about AI security has focused on models, when the durable problem has always been the data underneath them," said Rajesh Parthasarathy, Founder and Chief Executive Officer at Mage Data. "We have spent several years protecting enterprise data wherever it lives — and AI does not change that discipline, it stretches it. The same policy that governs a database must now govern a training pipeline, a prompt, and an agent's answer. That is what we are delivering: not a new product for a new fear, but our platform extended to every point where data meets AI."

"Most enterprises have not chosen between AI adoption and data protection — they have discovered, after the fact, that the decision was made for them," said Anil Bhat, Senior Vice President and Chief Technology Officer at Mage Data. "Blocking AI outright simply relocates the risk to personal accounts and unmanaged devices. We engineered these capabilities around one principle: protect the data, not the tool. Whether data enters a training pipeline, a public assistant, or a vendor copilot, the same policy applies."

Availability

Data Security and Privacy for AI is available now. Book a demo to see the capabilities in action, then run a proof of concept (PoC) in your own environment.
For more information about Data Security and Privacy for AI, visit www.magedata.ai/products/data-security-privacy-for-ai

About Mage Data

Mage Data is a leader in data security and privacy software for global enterprises. Mage Data has been recognized as a Leader in Data Security Platforms by KuppingerCole and a Champion in Test Data Management by Bloor Research, and has been included as a notable vendor in Forrester's Data Security Platforms Landscape, Q3 2026. Its award-winning and patented platform helps organizations navigate privacy regulations while protecting sensitive data. Mage Data is also the highest-ranked vendor in Test Data Management and Data Masking in Gartner Peer Insights over the past 12 months, reflecting strong customer validation of its capabilities and impact. Mage Data's client roster includes Fortune 500 companies, leaders in the financial and healthcare sectors, and Ivy League universities.

For more information, visit www.magedata.ai

Media Contact
bards@magedata.ai | +1 212 203 4365

Rajesh Parthasarathy
Mage Data
+1 212-203-4365
email us here
Visit us on social media:
LinkedIn
Instagram
YouTube
X
