POINT OF VIEW

Roll Out AI Applications and Agents Faster with Policy, Spend, and Data Under Your Control



Executive Summary

Business leaders are worried about falling behind on AI, pushing for AI tools to be rolled out immediately to employees. To meet this demand, IT teams plan to connect applications to a public large language model (LLM) service or to host their own model.

Those projects create immense security risk, yet teams lack mature AI runtime security and the budget to build it from scratch. What's needed is a cost-effective AI gateway that applies runtime guardrails, compliance policies, and cost controls as teams build AI tools and skills. This way, teams can complete AI projects with confidence and contain the risk while the security program matures.

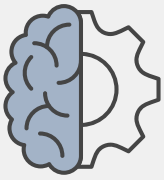

Nine in 10 organizations use AI regularly in at least one business function.[1]

Pressure to Adopt AI Quickly

AI adoption is now broad, and when business leaders see peers ship AI features, they ask their teams to move faster. This prompts teams to start experimenting with AI. For example, one team connects a customer application to a public LLM service. Another team runs an open-weight model on internal graphics processing units. Each project looks small, but together they create a large AI attack surface that current security controls do not inspect.

Security teams then face a choice between AI adoption speed and security. AI projects move faster than security reviews. If security holds a project until controls exist, the team adopts shadow AI without security review. If security approves a project without controls, the organization takes on risk it cannot measure.

Why Early AI Projects Need New Controls

Risk follows adoption. In 2026, more than 20% of organizations reported a breach that targeted AI models or apps.[2] While network firewalls and data loss prevention (DLP) tools inspect packets and payloads, they do not read the meaning of a prompt or a model response. AI projects face four risks that those tools miss:

- **Prompt injection:** Attackers hide instructions in prompts or in content an agent reads, and the model follows them.
- **Data leakage:** Users paste personally identifiable information (PII) into prompts, and models return confidential data in responses.
- **Runaway spend:** Inefficient prompts and attacks that flood models with costly requests drain a limited AI budget.
- **Missing evidence:** Auditors expect a record of who used which model and with what data.

A Safer Way to Build AI Tools and Skills

Safe AI adoption starts with one control point between the applications and models. The control point reads AI traffic in context. It inspects each prompt before it reaches a model, all tool calls, and each response before it reaches a user. It applies one set of policies to each model that teams try. It also logs each exchange, so teams have evidence ready when auditors ask. With that control in place, security teams can approve new AI projects with clear limits on risk and spend.



46.9% of enterprises report AI spend over budget.[3]

How to Protect AI Projects

The self-hosted AI gateway sits between the models and the AI apps and agents that use them. This way, it can inspect each prompt before it reaches a model and each response before it reaches a user. The gateway also inspects an agent's tool responses and tool calls before the agent acts on them, including calls to Model Context Protocol (MCP) tools. The gateway:

- Blocks prompt injection, jailbreaks, toxic output, and harmful MCP tool calls at runtime
- Detects PII and other sensitive data in prompts and responses, then redacts or blocks it
- Contains each AI project with workspaces, roles, and model access controls for each user, team, and project
- Caps AI spend by user, team, project, and model, and routes routine requests to lower-cost models
- Adds more than 50 predefined compliance policies in one click and logs each prompt, output, and tool call for auditors
- Keeps prompts, data, and logs inside the organization's environment

Start with One Project, Then Expand

When applications and agents connect through one OpenAI-compatible application programming interface, the first project needs no re-architecture. Teams start with one application or agent and run a proof of concept to measure guardrail accuracy, latency, and cost visibility. They extend the same policies to each new AI project. This enables organizations to adopt AI at the pace the business needs, with guardrails in place from the first prompt.

One Control Point for AI Security, Policy, and Spend

Organizations do not have to choose between AI adoption speed and security. Traditional firewalls and DLP tools fall short because they cannot interpret the context or meaning of AI traffic. That gap leaves enterprises exposed to prompt injections, data leakage, runaway budgets, and compliance failures.

A self-hosted AI gateway gives security teams one control point that inspects each prompt and response at runtime. With runtime guardrails in place, teams can deploy AI tools quickly and keep control over risk, spend, and data.

Learn more

[1] The State of AI, McKinsey & Company, August 25, 2026. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai#

[2] Cost of a Data Breach Report 2026, IBM, July 29, 2026. https://www.ibm.com/reports/data-breach

[3] 46.9% of Enterprises Report AI Spend Over Budget in 2H 2026, Futurum, September 10, 2026. https://futurumgroup.com/press-release/46-9-of-enterprises-report-ai-spend-over-budget-in-2h-2026/

FORTINET

www.fortinet.com

Copyright © 2026 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's SVP Legal and above, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.

October 20, 2026 3:11 PM / MKTG-4439-0-0-EN