



Rethink your data perimeter: Secure generative AI from the inside out

As public sector organizations scale generative artificial intelligence (AI), many find their existing data protection strategies must transform fast.

Traditional data loss prevention (DLP) focuses on securing the system perimeter, but generative AI operates within the system, drawing on sensitive information scattered across hybrid environments. Without rethinking DLP for generative AI, organizations risk exposing sensitive information.

[Forcepoint](#), built on Amazon Web Services (AWS), addresses this challenge through its AI-powered data security platform. Forcepoint's AI mesh technology automatically discovers, classifies, and protects sensitive information across multiple environments, providing granular controls specifically designed for generative AI workflows—helping organizations scale generative AI while preventing sensitive data loss to large language models (LLMs).

Why Forcepoint?



Mission-aware, AI-powered data protection

With over 1,700 out-of-the-box classifiers and templates for 80+ countries and industries like healthcare and finance, Forcepoint enables contextually aware, reusable, and adaptive DLP policies for generative AI use. Internal departments can define their own data rules and tune protections as compliance needs evolve.



Zero trust for internal AI workflows

Forcepoint moves beyond traditional perimeter-based DLP to monitor and control internal workflows with zero trust protocols, preventing data exposure within generative AI applications and multi-AI-agent workflows. Forcepoint's risk-adaptive data protection evolves based on an individual user's risk level.



Unified DLP management across data channels

Create data security policies once and apply them to web traffic, SaaS applications, email, and endpoints. With real-time audit tools, security teams can quickly remediate generative AI risks fast.



Operational scale with built-in data privacy

Forcepoint automates classification and risk assessment across fragmented legacy, hybrid, and multi-cloud environments with no data centralization required. Data is never moved or stored from its original location, obeying strict privacy and sovereignty standards.

Real-world impact: Blocking data breaches before they happen

A public sector agency recently integrated an enterprise-level commercial LLM to support staff productivity. But with AI tools increasingly woven into daily workflows, leadership was concerned about the risk of sensitive data leaking through prompts.

Using Forcepoint's AI mesh and DLP capabilities, the agency was able to automatically identify and block a user's attempt to upload personally identifiable information into the model, preventing a potential data breach.

"Just stopping one data breach can represent huge ROI, because even one data leak could represent very significant damage," said Ronan Murphy, Chief Data Strategy Officer at Forcepoint. In fact, according to the [2024 Cost of a Data Breach report](#) by IBM, the global average cost of a data breach is over \$4 million.

"AWS offers a huge degree of diligence and security controls in how we build and provide service for customers."

— Ronan Murphy, Chief Data Strategy Officer, Forcepoint

Is Forcepoint right for my organization?

Ask yourself these key questions to determine if your teams would benefit from Forcepoint's AI-powered DLP platform:

- 1 Do you understand where all your data lives?** Public sector data estates often span decades-old systems, hybrid-cloud platforms, and file shares. If you can't map where sensitive information exists across your environments, you likely have blind spots that could lead to data exposure.
- 2 Have you classified data by mission, department, and regulatory impact?** Every department handles data differently. Human resources, finance, public health, and transportation teams all have unique sensitivities. Without granular, mission-aligned classification reflecting data contexts, you cannot apply effective controls for generative AI.
- 3 Can you define what data your AI systems are allowed to access?** If you can't clearly describe which datasets each AI tool can use, or what each user can access, your AI initiatives may expose sensitive information.
- 4 Is your DLP strategy built for internal AI workflows?** Legacy DLP was designed to stop data from leaving your network, but AI now lives inside the firewall. Modern protection means monitoring collaboration tools and internal agents before an AI prompt becomes a data leak.

"With generative AI, data really has become the new oil. It allows organizations to generate incredible operational efficiencies, insights, and innovation. But on the other side of that coin, it also presents very significant risk."

— Ronan Murphy, Chief Data Strategy Officer, Forcepoint

How AWS powers Forcepoint's platform

Cloud-native and secure by design

Built entirely on AWS, Forcepoint uses in-place data scanning and lightweight collectors to meet strict public sector privacy and residency requirements. Using AWS's best-in-class security features, Forcepoint helps organizations maintain compliance with NIST, CMMC, FIPS, and other frameworks.

Enhanced protection for generative AI applications

Forcepoint integrates with AWS services like Amazon Bedrock, Amazon SageMaker, and Amazon Q to enforce DLP policies across entire generative AI workflows. This includes scenarios like Retrieval-Augmented Generation (RAG), where sensitive data could be exposed through prompts or knowledge bases.

Procurement-ready for the public sector

Available in AWS Marketplace and eligible under AWS Enterprise Discount Program (EDP) commitments, Forcepoint supports streamlined procurement for budget-constrained teams navigating complex approval cycles.



Explore how Forcepoint helps public sector organizations protect sensitive data while embracing generative AI.



Learn more →