



TRACK 2

Prevent data loss with generative AI

As public sector organizations scale generative artificial intelligence (AI) from pilot to production, they must rethink how they approach data security. Traditional data loss prevention (DLP) strategies, which create perimeter-based controls, are not designed for the way generative and agentic AI tools operate: inside the system and across internal workflows.

Some organizations respond by taking an overly restrictive approach, blocking generative AI access broadly to reduce risk. But this kind of blanket policy often backfires by frustrating users and driving unsecure workarounds. A more effective path is to adopt flexible, granular DLP policies that can be tailored to different roles, data types, and use cases to protect sensitive information without slowing progress.

This track shares real examples from Amazon Web Services (AWS) and AWS Partners to help public sector leaders prevent sensitive data loss and maintain trust, agility, and effectiveness in their generative AI applications

Why choose AWS Partners?

Built-in data protection for generative AI solutions

AWS Partners help public sector teams modernize their DLP strategies with comprehensive generative AI controls and visibility. These solutions are built on the security and compliance foundations of the AWS Cloud and integrate with services like Amazon Bedrock and Amazon SageMaker to provide data protection across the full generative AI lifecycle.

Explore curated generative AI DLP solutions in the [Partner Expo](#).

Why public sector organizations need targeted DLP solutions for generative AI

To ensure generative AI deployments scale securely across organizational-wide use, public sector organizations can use modern DLP tools to:



Prevent exposure from shadow AI

Even in highly regulated environments, employees may use unsanctioned tools to increase their productivity and efficiency. Generative AI DLP solutions provide visibility into where and how generative AI is being used, helping organizations address risks without stifling innovation.



Balance access and control

Rather than blocking generative AI entirely, nuanced DLP policies enable safe, productive use. Users can explore AI tools while restricting sensitive data uploads, downloads, or other unsanctioned content sharing.



Scale securely across environments

Generative AI often interacts with data stored across cloud, on-premises, and hybrid systems. Modern AI-aware DLP policies that apply consistently across these enterprise-wide environments are critical for maintaining compliance and operational integrity as usage expands.



Build trust in generative AI systems

Clear, consistent guardrails help make sure that generative AI adoption doesn't come at the expense of data privacy or compliance. Generative AI DLP strategies let leaders confidently scale new use cases while maintaining public trust and integrity.

How AWS provides foundational security

AWS provides the foundation for DLP solutions that help public sector organizations scale generative AI safely. With cloud-native controls and integrations across generative AI services, AWS helps organizations:

- ✓ **Apply centralized security and DLP policies** across the all generative and agentic AI behavior, including input prompts, model outputs, and multi-step workflows like retrieval-augmented generation (RAG), with services like [Amazon Bedrock Guardrails](#).
- ✓ **Maintain data sovereignty and privacy** using in-place data storage and scanning, without moving sensitive information.
- ✓ **Meet public sector compliance requirements** with FedRAMP-authorized infrastructure and support for frameworks like NIST, CMMC, and FIPS.

How to prevent data loss in your generative AI solutions

1

Begin with visibility into data dependencies and generative AI usage

Before launching internal generative AI initiatives, public sector organizations need to determine two things: where their sensitive data lives and how generative AI is already being used by employees in the organization. First, conduct a data-centric risk assessment by mapping how your data flows across cloud, on-premises, and hybrid environments.

Next, assess generative AI activity across your organization. Which tools are in use, both officially or unofficially? Who is using them, and what kind of data are they sharing? Without visibility, sensitive information can easily be entered into unsecured large language models (LLMs), creating irreversible exposure. This groundwork helps organizations mitigate the risk of shadow AI and establish clear governance from the start.

2

Customize generative AI DLP controls based on users and use case

Not all data is equally sensitive, and not all teams use AI the same way. Effective DLP must be flexible enough to reflect departmental needs, role-based access, and the specific use case for AI tools. In fact, overly broad restrictions can send employees to find unsecure workarounds and shadow AI tools in secret. Instead of blocking generative AI outright, consider policy-based guardrails like browser isolation, prompt-level filtering, and upload/download restrictions to manage risk without limiting productivity.

3

Educate your teams with a new DLP mindset shift

Traditional DLP focuses on stopping data from leaving the infrastructure perimeter. But generative AI changes the equation: tools like AI co-pilots, internal LLMs, and agentic AI now operate *inside* the perimeter. Public sector organizations need to shift how they approach security to apply zero trust controls within systems and applications to prevent data loss via generative AI and agentic AI.


But technology alone isn't enough. As you scale AI, build user education into your rollout to help your teams understand not just what behaviors to change, but why secure generative AI practices matter to your organization's broader mission and risk posture.



Explore featured solutions from AWS Partners

AWS Partners are already helping public sector organizations reduce data loss risks while embracing generative and agentic AI innovation.

Visit the “Partner Expo” to explore how these
solutions can support your organization’s goals →



KEY CONTRIBUTORS

Mehmet Bakkaloglu | *Principal Solutions Architect for ISV Partners, AWS*

Kunal Sharma | *Senior Solutions Architect, AWS*

Gartner Press Release, Gartner Predicts 40% of AI Data Breaches Will Arise from Cross-Border GenAI Misuse by 2027, February 17, 2025

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.