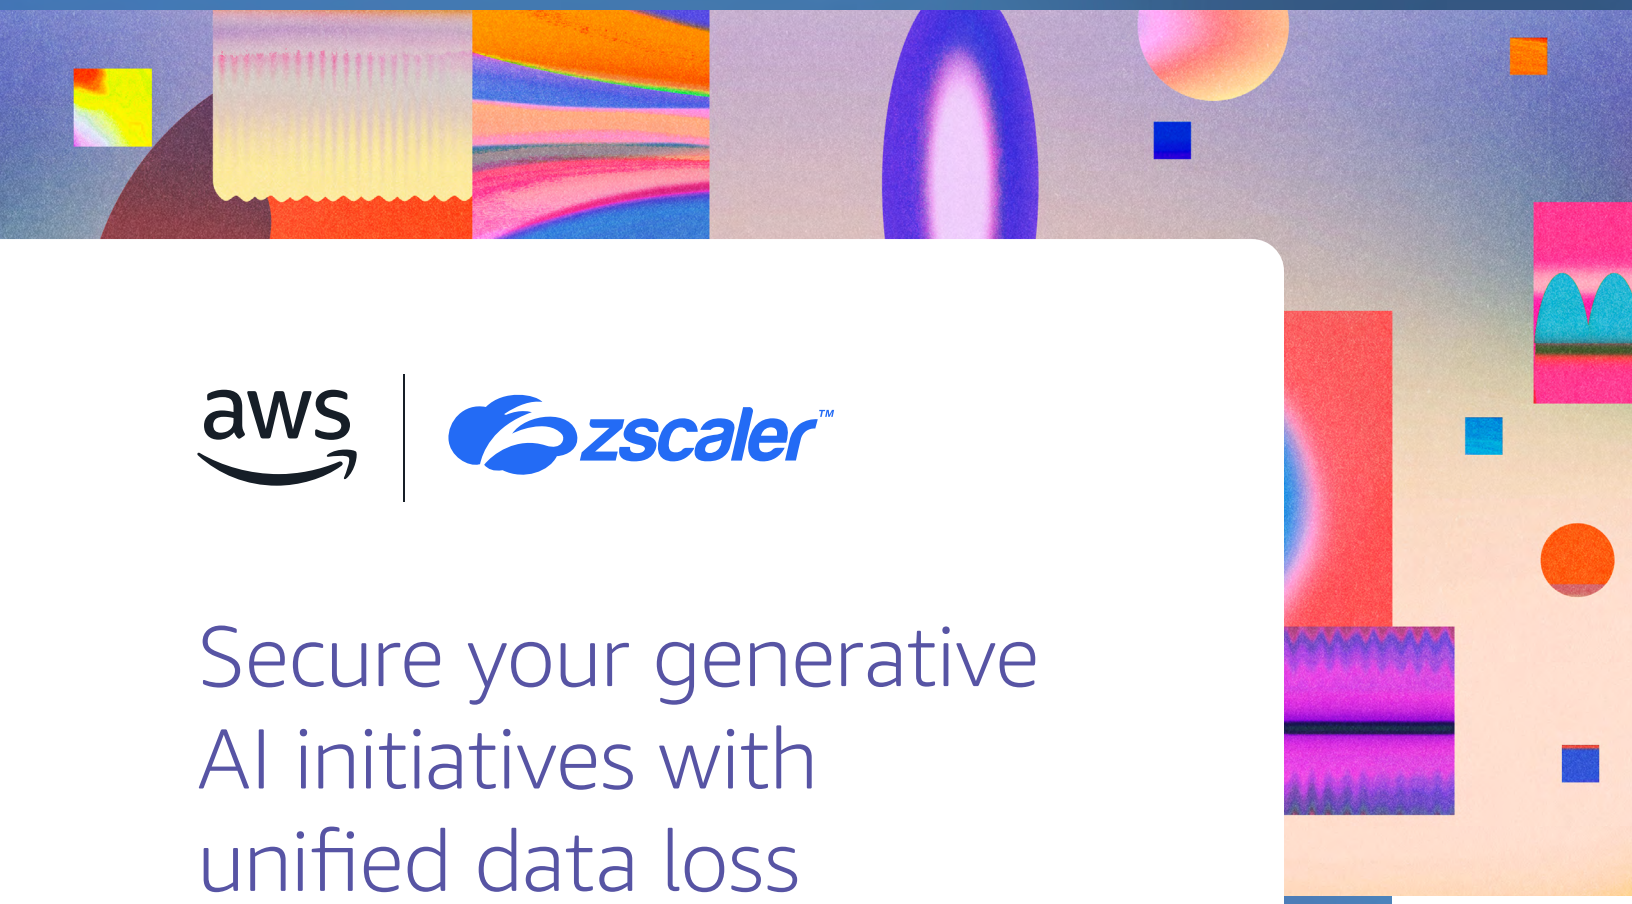

Secure your generative AI initiatives with unified data loss prevention

As public sector organizations scale generative artificial intelligence (AI) from pilot to production, one risk rises sharply: unintended data exposure.

As public sector organizations scale generative artificial intelligence (AI) from pilot to production, one risk rises sharply: unintended data exposure. Gartner® predicts that "by 2027, more than 40 percent of AI-related data breaches will be caused by improper use of generative AI across borders."

Zscaler helps public sector organizations secure these AI borders. The Zscaler platform integrates with Amazon Web Services (AWS) to embed data loss prevention (DLP) processes directly into an organization's entire IT environment, including granularly across generative AI applications. With prompt-level visibility, consistent DLP policy enforcement, automated security incident response, and more, organizations can use Zscaler to accelerate generative AI adoption across their teams securely and confidently.

How Zscaler proactively secures data for generative AI

Real-time security enforcement to AI

Zscaler applies AI-aware security policies directly in the network path, so controls happen in real time without slowing users down.

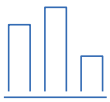

Granular visibility into AI usage

Understand exactly which AI tools are being used, who is using them, and what prompts or data are being submitted, even across commercial AI apps.

Built-in browser isolation

Isolate risky activity automatically without outright blocking all access. Users can view content safely, but uploads, downloads, and copy/paste actions are restricted.

Unified zero trust architecture

All protections run through a single platform, so you can extend your current security policies and compliance needs to all AI use cases.

Faster response to reduce security risks

Automatically flag, route, and resolve AI-related incidents across your whole organization, helping teams protect sensitive data without adding manual workload.

Built for public sector needs

Zscaler's FedRAMP-authorized platform is trusted by more than a dozen U.S. federal cabinet agencies.

zscaler™ *at a glance*

2.19M annual average customer savings *due to increased staff productivity, security risk avoidance, and more*

27% faster resolution *of actual data loss instances*

14% average reduction *of major data-related security incidents*

37% faster scaling *to application deployment*

33% faster completion *of compliance reports*

5 steps to prevent data loss in generative AI

Generative AI offers major productivity gains, but without the right guardrails, the risks regarding data loss are real. Zscaler recommends a practical, policy-driven approach to help public sector organizations adopt generative AI at scale securely.

1 **Start with visibility.** To protect your data, you need to know what AI tools are in use across your organization, who is using them, and what kind of data is being shared with them.

2 **Avoid the "block everything" trap,** which can drive employees to seek out shadow AI tools to ease their workload. Instead, consider nuanced, targeted DLP solutions like browser isolation to make certain AI tools available, but restrict the data users can upload, prompt, receive, download, and more.

3 **Apply centralized DLP policies** to maintain consistent security standards across all your systems, including generative AI applications. This protection is particularly important for multi-agent workflows where sensitive data passes through interconnected AI agents, creating more points where data needs to be secured.

4 **Educate users as part of your security culture.** Many security incidents stem from a lack of awareness about the dangers of data loss. Build messaging and end user training strategies into your generative AI rollout plan to make sure your teams understand the importance of DLP.

5 **Automate your incident response.** As the scale of security alerts grows, manual triage isn't sustainable. You can use generative AI tools like Zscaler's Workflow Automation to automatically identify risks, initiate the right escalation path, take action as allowed, and close out tickets to speed IT efficiency.

Is Zscaler the right fit for my organization?

Ask yourself these key questions to determine if your teams would benefit from Zscaler's DLP tools:

- *Do we know which generative AI tools are being used across our environment? Do we know exactly what staff is exchanging with LLMs?*
- *Are our existing DLP policies consistently applied to generative AI use cases?*
- *Is our AI IT strategy flexible? Can we adjust our security controls to address new AI-related risks as they emerge?*
- *Can we manage and resolve AI incidents at scale without overwhelming our security teams?*

Explore how Zscaler helps public sector organizations adopt generative AI with visibility, control, and confidence.

Learn more about securing generative AI with Zscaler or contact Zscaler today. →

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