

AI, Data Security, & Data Loss Prevention

Protecting Sensitive Data in the Age of Generative AI

Presented by: Jarrod Daake

October 21, 2026

Agenda

01 The AI Adoption Wave

02 The New Risk Landscape

03 DLP Fundamentals

04 Best Practices

05 A 90-Day Roadmap

06 Q & A



The AI Adoption Wave



DATAVIZION

AI Use Has Tripled

45%

of employees are now regular AI tool users, up from 15% a year earlier

67%

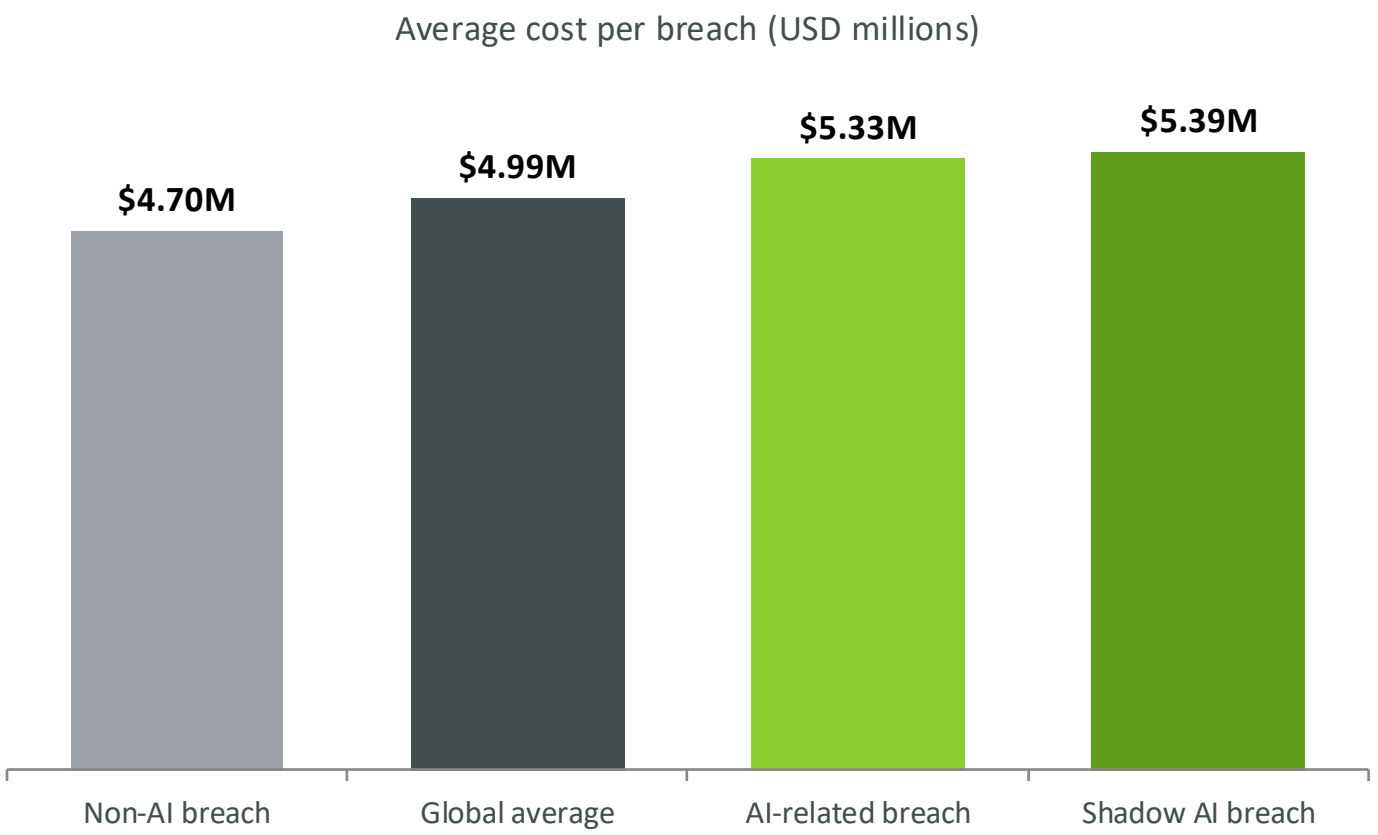
of AI users on corporate devices sign in with non-corporate accounts

~40%

of employee AI interactions involve sensitive data: prompts, pastes, and file uploads

DATAVIZION

Breaches Cost More, and AI Adds a Premium



$4.99M

Global average breach cost, a record high and up 12% year over year

$11.5M

Average breach cost in the United States

DATAVIZION

The New Risk Landscape



DATAVIZION

Six New Paths for Data Loss

Shadow AI

Unapproved tools and personal accounts with no logging, retention controls, or contract

Prompt leakage

Employees paste customer data, code, or financials into public chatbots

Oversharing

AI assistants surface files users can technically access but never should see

Prompt injection

Hidden instructions in content trick AI agents into exfiltrating data

Over-privileged agents

AI agents and service accounts with broad, unmonitored access to systems

AI-powered attacks

Deepfake impersonation and AI-written phishing that bypass human checks

The Governance Gap

92%

of organizations with an AI-related breach had no proper AI access controls

43%

of breached organizations reported a shadow AI incident, up from 20%

32%

of breached organizations have an AI governance policy in place

The problem usually isn't AI itself. It's AI deployed without policy, identity controls, or visibility.

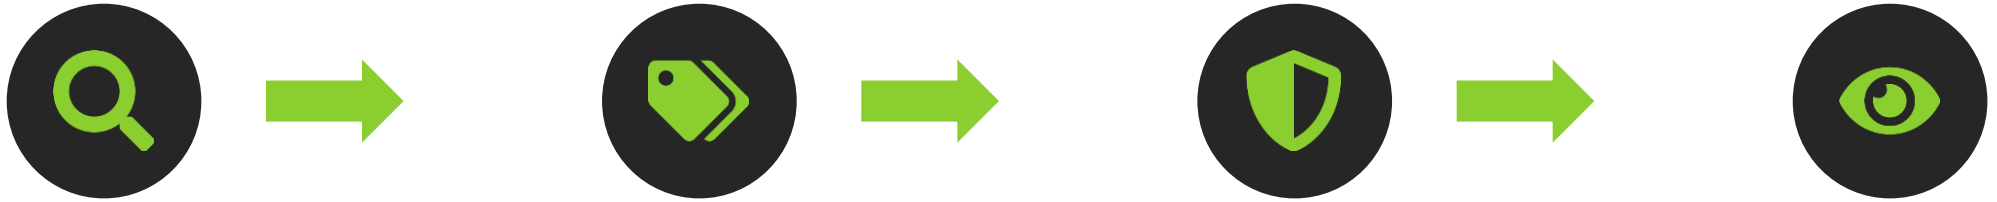
DATAVIZION

DLP Fundamentals



DATAVIZION

How Data Loss Prevention Works



Discover

Find where sensitive data lives across endpoints, email, cloud apps, and file shares

Classify

Label by sensitivity: public, internal, confidential, regulated

Protect

Apply policies that block, warn, encrypt, or restrict sharing by label

Monitor

Log activity, alert on risky behavior, and tune policies over time

Data at rest

Stored files, databases, SaaS

Data in motion

Email, web uploads, AI prompts

Data in use

Copy/paste, print, screen, apps

DATAVIZION

Best Practices



DATAVIZION

Govern AI Before You Scale It



Policy is the foundation. Every technical control enforces a decision made here.

1. **Publish an AI acceptable use policy**
   Name approved tools, the data that must never go into them, and who to ask

2. **Provide sanctioned enterprise AI**
   A safe option with business accounts, no training on your data, and audit logs

3. **Classify data and tie rules to labels**
   "Confidential" and "Regulated" data gets blocked from unapproved AI destinations

4. **Vet AI vendors and features**
   Review data retention, model training, residency, and subprocessors before approval

5. **Assign ownership**
   Name an AI governance owner and make security and AI teams work from one plan

DATAVIZION

Layer Your Technical Controls



Inspect AI traffic

Use browser or network DLP to scan prompts and uploads to AI sites; block or warn on sensitive content



Least-privilege identity

Clean up file permissions before rolling out AI assistants; scope agent and service-account access



Encrypt & label

Persistent labels and encryption keep protection attached to the file wherever it goes



Control unsanctioned apps

Discover AI apps in use and block personal-account sign-ins on corporate devices



Endpoint, email & cloud DLP

Cover copy/paste, USB, print, email, and SaaS sharing with one consistent policy set



Log & alert

Centralize AI and DLP logs in your SIEM; alert on bulk uploads and unusual access

DATAVIZION

Make People Part of the Control

Train on real scenarios

Show what's OK to paste and what isn't, using examples from your own business

Coach in the moment

Use DLP warnings that explain why and offer the approved alternative, not just a block

Plan for AI incidents

Add AI data exposure, deepfake fraud, and agent misuse to the incident response plan

Review and improve

Revisit policies quarterly as tools, regulations, and usage patterns change

Employees just want to get work done, so make the safe path the easy path.

DATAVIZION

A 90-Day Roadmap



DATAVIZION

A Practical 90-Day Roadmap

1 — 2 — 3

Days 1–30

Assess

- Inventory AI tools in use, sanctioned and not
- Identify your most sensitive data types
- Publish an interim AI use policy

Days 31–60

Enable & protect

- Roll out an approved enterprise AI tool
- Fix overshared files and permissions
- Turn on DLP in audit mode for AI sites

Days 61–90

Enforce & measure

- Move DLP rules from audit to warn/block
- Launch scenario-based staff training
- Report metrics and set a quarterly review

DATAVIZION

Key Takeaways

Assume AI is already in use
Get visibility first, then shape behavior with a sanctioned option.

Governance drives security
Policy, ownership, and classification make every control enforceable.

Layer your defenses
DLP across prompts, endpoints, email, and cloud.

AI can defend, too
Extensive security AI and automation cut breach costs by $1.93M and containment time by 65 days.

DATAVIZION

Q & A



DATAVIZION