AI Privacy & Data Governance

The Ducats Group

What your employees are unknowingly sharing with AI tools, why it matters more than most organizations realize, and the practical governance steps that close the gap before it becomes a liability.

> *"Most AI data risks in organizations aren't malicious — they're well-intentioned employees trying to do their jobs faster, unaware that the tool they're using is storing, reviewing, or training on everything they type."*

THE SIX MOST COMMON AI DATA LEAKAGE SCENARIOS

CRITICAL RISK

Pasting client data into consumer AI

An employee pastes a client contract, financial statement, or customer list into ChatGPT Free or a personal Gemini account to get a summary or analysis. That data may now be stored by the AI provider and potentially used for model training. The client never consented. Your NDA may have just been breached.

Fix: Mandate business/enterprise tiers only for work tasks. Block consumer AI URLs at the network level if needed.

CRITICAL RISK

Sharing PII in AI prompts

Employee names, Social Security numbers, medical information, salary data, and customer contact details are routinely included in prompts without a second thought. Under GDPR, CCPA, and HIPAA, transmitting PII to a third-party AI service without a Data Processing Agreement may constitute a reportable data breach.

Fix: Train staff to anonymize data before prompting. Establish a clear rule: if it would be sensitive in an email, it's sensitive in a prompt.

HIGH RISK

Browser AI extensions with broad permissions

Many popular AI browser plugins — writing assistants, email helpers, meeting summarizers — request permission to read all page content. This means they can silently read every document, email, and web page the employee visits, including authenticated internal systems.

Fix: Establish an approved browser extension policy. Review and restrict extension permissions via MDM or browser management.

HIGH RISK

AI meeting recorders without consent

Tools like Otter.ai, Fireflies, and similar apps join calls and record everything — including confidential strategy discussions, M&A conversations, and personnel matters. Many jurisdictions require explicit consent from all parties. Uninvited bots on client calls may violate wiretapping laws.

Fix: Define which recording tools are approved. Require explicit consent disclosures. Block unapproved bots from joining company meetings.

MODERATE RISK

Proprietary IP in code assistants

Developers using GitHub Copilot (free tier), ChatGPT, or similar tools to write or review code may inadvertently submit proprietary algorithms, API keys, database schemas, or internal architecture details. Some free-tier code assistants retain submitted code for model improvement.

Fix: Use enterprise code assistant tiers with explicit no-training guarantees. Add secrets scanning to CI/CD pipelines to catch exposed credentials.

MODERATE RISK

AI-generated content with hidden IP risk

AI tools trained on copyrighted material may reproduce protected text, code, or images in their outputs. Organizations publishing or commercializing AI-generated content without review may unknowingly infringe on third-party intellectual property — and bear the legal liability.

Fix: Establish a human review requirement for all externally published AI-generated content. Use IP-indemnified tiers where available (e.g., Microsoft Copilot Commercial).

Building Your AI Governance Framework

The Ducats Group

Platform-by-platform privacy facts, the regulatory landscape, and a practical policy checklist your organization can implement today — no legal team required to get started.

PLATFORM PRIVACY QUICK REFERENCE — WHAT ACTUALLY HAPPENS TO YOUR DATA

Platform / Tier	Trains on your data?	Data stays in your tenant?	DPA available?	Best for business use?
ChatGPT Free	Yes — by default	No	No	No — personal use only
ChatGPT Plus	Opt-out available	No	No	Caution — opt out first
ChatGPT Team/Enterprise	No	OpenAI servers	Yes	Yes
MS Copilot (M365)	No	Your M365 tenant	Yes (MSFT)	Yes — best for M365 orgs
Claude Free/Pro	May use for safety	No	No	Caution for sensitive data
Claude Team/Enterprise	No	Anthropic servers	Yes	Yes
Gemini (personal)	Yes — by default	No	No	No — personal use only
Gemini for Workspace	No	Your Google tenant	Yes (Google)	Yes — best for GWS orgs

REGULATORY LANDSCAPE — WHAT APPLIES TO YOUR ORGANIZATION

EU / UK

GDPR

GENERAL DATA PROTECTION REGULATION

Any EU/UK personal data sent to an AI tool requires a lawful basis and a Data Processing Agreement with the provider. Violations carry fines up to 4% of global annual revenue.

US — HEALTHCARE

HIPAA

HEALTH INSURANCE PORTABILITY ACT

Protected Health Information (PHI) cannot be shared with AI tools unless a Business Associate Agreement (BAA) is in place. Most consumer AI tools do not offer BAAs.

US — CALIFORNIA

CCPA / CPRA

CALIFORNIA CONSUMER PRIVACY ACT

California residents have rights over their personal data. Sharing customer PII with AI tools without disclosure in your privacy policy may constitute a violation — even for non-California companies.

US — FINANCE

GLBA / SOX

FINANCIAL DATA REGULATIONS

Financial institutions must protect non-public customer information. Using unvetted AI tools to process financial records may conflict with GLBA safeguard rules and SOX data integrity requirements.

YOUR AI GOVERNANCE CHECKLIST — START HERE

- ☐ **Approved AI tools list** published and communicated to all staff
- ☐ **AI Acceptable Use Policy** written, signed off, and distributed
- ☐ **Data classification policy** defines what can/cannot enter AI prompts
- ☐ **DPAs / BAAs in place** with every AI vendor processing personal data
- ☐ **Consumer AI blocked or restricted** on corporate devices and networks
- ☐ **Privacy policy updated** to disclose AI tool usage to customers
- ☐ **Staff trained** on what data types are prohibited in AI prompts
- ☐ **Browser extension audit** completed — unapproved AI plugins removed
- ☐ **AI meeting recorder policy** defined with consent requirements
- ☐ **AI output review process** for any externally published content

REALITY CHECK Most organizations are 6–18 months behind on AI governance. The risk isn't theoretical — regulators in the EU, UK, and US are actively issuing guidance and beginning enforcement actions related to AI data handling. **The organizations that act now are building a defensible position. Those that wait are accumulating liability.** The good news: a basic governance framework can be stood up in 30 days with the right focus.

TDG TAKE **The Ducats Group perspective:** AI governance isn't a legal department problem — it's an operational one. The policies, tool approvals, and training your team needs are straightforward to build. What's missing in most organizations isn't knowledge of what to do — it's ownership of who does it. Designating an AI governance owner and giving them a 90-day mandate to close these gaps is the highest-leverage move most mid-market organizations can make right now.