GUIDE

CONCENTRIC Ai

The 7 Deadly Sins of GenAI Security

and how to redeem yourself before it's too late

Welcome to the Confessional

Enterprises are tripping over themselves to embrace GenAI. But along the way, they're also committing repeatable, predictable, egregious and downright embarrassing security sins.

Most CISOs already know the risks, but in the chaos of adoption, it's easy to slip into bad habits that trade security for speed.

The price they're paying for these sins: sensitive data exposure, shadow GenAI everywhere, and a governance strategy that looks great on paper but fails in practice.



GenAI brings new challenges while it exposes weaknesses you already had. These "sins" are likely alive and kicking in your environment right now.

The Seven Sins Unveiled

GenAI has turned enterprise security into a mirror, and every organization has some unflattering reflections going on. Ignoring these sins is an even bigger sin in itself, so let's learn about what they are and why they're wreaking havoc for security teams and CISOs.



Prompt Gluttony
Employees paste entire contracts, financial forecasts, or source code into ChatGPT because it feels faster. That one "quick paste" can be a compliance nightmare.



Permission Sloth
Over-permissioned drives and folders left wide open across M365 or Google Workspace. Copilot and Gemini happily surface them to anyone who asks.



Classification Envy
Relying on regex or filenames to catch sensitive data in an era when semantic search can see right through your shortcuts.



Context Pride
Believing AI tools understand the difference between a press release and an NDA draft. No spoiler alert required for this one, you know the score.



Shadow GenAI Lust
Sneaking off with unapproved public tools because governance feels too slow, too rigid, or too focused on saying "no."



Output Greed
Copy-pasting AI-generated text into decks, emails, or code without realizing it may contain sensitive context that spreads downstream.



Governance Wrath
Swinging the pendulum too far with draconian policies. The stricter the controls, the faster employees look for ways around them and begin adopting shadow AI.

These aren't hypotheticals. They're the daily ways GenAI is already exploiting cracks in your defenses.

Redemption Roadmap

Sins can be forgiven but only if you act quickly. The path to progress shall not be paved with chasing perfection, but with visibility, balance, and smarter protections aplenty that actually lead your teams back to the truth.

Governance done right is less about wrath and more about the right path forward.

Actionable Guidance

Confess early

1

Absolve your organization by acknowledging where GenAI is already in use, including shadow GenAI.

Classify with context

2

Bring true order to your data, by moving beyond regex and applying semantic labeling that understands meaning and context.

Tighten access without strangling productivity

3

Least privilege is essential if you want Copilot to be safe. Done right, it's not restrictive but a huge part of good governance.

Sanitize outputs

4

Don't let leaks spread downstream. Treat AI-generated content as carefully as the inputs.

Govern by design, not wrath

5

Build flexible policies that adapt as usage evolves, so you stay ahead of the next wave, and not engulfed by it.

Pro tip

The worst mistake isn't gluttony, sloth or greed. It's pretending you have nothing to confess.

Start turning your GenAI sins into strengths—before regulators, auditors, or your own employees call you out.

Book a Demo