GTK CYBER

EXECUTIVE RESOURCE

The CISO AI Risk Checklist

The AI exposure that lands on a security team this quarter is already inside the building, not a future nation-state scenario. This is the posture that addresses the risk that is actually here.

1. Inventory the AI you already have

You cannot govern what you cannot see. Start here, before policy.

- [ ] Procurement records: any contract mentioning AI, ML, or LLMs
- [ ] Vendor tools with newly added AI features (CRM, ITSM, security products, productivity suites)
- [ ] Developer tooling: Copilot, Cursor, internal API keys to model providers
- [ ] Most organizations are surprised by what shows up in the developer-tooling category

2. Treat shadow AI as a data-classification problem

- [ ] Define what data may and may not enter external AI systems
- [ ] Extend existing data classification and DLP controls (Microsoft Purview, Netskope, Zscaler)
- [ ] Do not treat this as an awareness-training problem; it is a technical-controls problem

3. Apply least privilege to AI agents

- [ ] Scope each agent's tools and credentials to only what its task requires
- [ ] Audit agent tool calls
- [ ] A document-summarizing agent does not need email-send; a support agent does not need database write

4. Ask AI vendors the hard questions

- [ ] What mitigations have you implemented against prompt injection?
- [ ] Can you demonstrate resilience against adversarial inputs?
- [ ] Is your model isolated from other customers' data at inference time?

5. Pick a governance framework and apply it

- [ ] NIST AI RMF as the operational framework
- [ ] EU AI Act as a compliance overlay if you have EU operations or customers
- [ ] ISO 42001 as the emerging certifiable management-system standard

6. Close the organizational gap

- [ ] Establish a cross-functional AI governance structure with authority to gate new deployments
- [ ] Build the technical literacy to evaluate AI-specific risk, not treat it as a compliance checkbox
- [ ] Name an owner for the intersection of security, legal, and business-unit AI deployments

Want this walked through with your team?

A one-day executive briefing, scoped to your AI risk. No technical prerequisites.

gtkcyber.com/executive-briefing

info@gtkcyber.com · 251-GTK-CYBER

Adapted from GTK Cyber practitioner analysis. Full article: gtkcyber.com/blog/what-cisos-get-wrong-about-ai-risk