AI Security Readiness Checklist



Presented to

Security Leaders

Presented by

Monixity Team

MONIXITY
Securing Today. Ensuring Tomorrow.

Who This Checklist Is For, What You'll Get, and How to Use It



This checklist is designed for CISOs and security leaders in critical, high-risk sectors such as healthcare, pharma, banking, utilities, and public services. These leaders are being asked to deliver AI value while protecting crown-jewel data and meeting strict regulatory expectations. Each step in this guide offers practical, security-first actions tailored to complex enterprise environments.

Readers will gain a clear view of the essential security controls for AI systems and copilots, from data scoping and least-privilege access to guardrails, logging, and governance. By working through the seven steps, organizations can establish a solid foundation for their AI security program, enhance operational resilience, and better safeguard sensitive information. The checklist is designed to help security and technology leaders make confident decisions as they scale AI across the business.

This checklist works best as a working session tool. Bring your security, data, and AI owners together, walk through each step, and score your current state. Use the results to agree on the top three actions you will take in the next 30 days to improve AI security.

AI Security Readiness Checklist: Steps 1 and 2

Laying the foundation with data scoping and least-privilege access



The first two steps of the 7-step checklist focus on what your AI is allowed to see and how it is allowed to see it. Step 1 maps crown-jewel data and defines what is in scope for AI, and Step 2 enforces least-privilege access. Together, they give CISOs and security leaders a defensible foundation for every other AI control.

Step 1 – Map crown-jewel data and define AI scope.
Identify your highest-value, highest-risk data (PHI, PCI, trade secrets, exec communications). Decide what is in and out of scope for AI use cases and record it so security, data, and engineering teams all work from the same boundaries.

Step 2 – Enforce least-privilege access from day one.
Use dedicated service identities for AI and grant only the minimum access needed for approved use cases. Review these permissions with security and keep them auditable, so there are no "god-mode" AI accounts and any incident has a limited blast radius.

Successfully implementing these first two steps ensures you can clearly answer, "What data can this AI reach, and why?" a critical starting point for any serious AI security strategy.

AI Security Readiness Checklist: Steps 3 and 4

Putting guardrails on retrieval and content safety



Step 3 – Put retrieval and tools behind guardrails

Limit what your AI can call and where it can look. Use an allow-listed set of tools and data sources with clear schemas and timeouts. Scope retrieval by tenant, business unit, or project instead of "search everything," and log which sources and tools were used for each answer.

Step 4 – Add content safety and DLP checks

Route prompts and responses through a content-safety layer to catch jailbreaks, harmful content, and policy violations. Add DLP rules to detect secrets, identifiers, and regulated data in AI outputs. When something risky appears, block or redact it, log the event, and alert the right team.

By getting Steps 3 and 4 right, you control what your AI can reach and how it behaves, sharply reducing the risk of data leakage and unsafe responses.

AI Security Readiness Checklist: Steps 5 and 6

Logging, monitoring, and AI acceptable-use policy



Steps 5 and 6 focus on visibility and governance. Step 5 makes AI activity fully observable through logging and integration with your security stack. Step 6 turns that into clear, human-readable rules through an AI acceptable-use policy and ownership model. Together, they make AI both monitorable and manageable.

Step 5 – Log AI activity into your security tools.

Log prompts, responses, retrievals, and tool calls with user IDs, and send them to your SIEM so AI activity is searchable, auditable, and can trigger alerts like any other security signal.

Step 6 – Publish a simple AI acceptable-use policy.

Create a 1–2 page policy that lists approved AI tools, clear do/don't rules for using them with sensitive data, and names owners for the AI product, the data, and security/compliance.

Implementing Steps 5 and 6 significantly strengthens your AI security posture by making AI activity observable, and employee behavior guided by clear rules and ownership.

AI Security Readiness Checklist: Step 7 and Beyond

Start small, test hard, then scale AI safely



Step 7 focuses on how you launch and grow AI safely. Instead of a big-bang rollout, you pilot with a limited group, test aggressively, and refine guardrails before scaling. This gives CISOs and security leaders a controlled path to expand AI without losing visibility or control.

Pilot with intention. Start with a small set of users and one or two clearly defined use cases, so you can observe real behavior and risk patterns.
Test hard before scaling. Run adversarial and red-team scenarios against your AI system, trying prompt injection, data exfiltration attempts, and policy violations. Use findings to adjust prompts, access, tools, and monitoring.
Scale with guardrails. Only expand access once logging, policies, and controls from Steps 1–6 are in place and you have clear criteria for success and rollback.

By treating Step 7 as an ongoing cycle of piloting, testing, and refinement, you ensure AI adoption stays secure, measurable, and aligned with your risk appetite as it grows.

Trusted by Organizations Where AI and Data Security Really Matter

Monixity supports security and data leaders in healthcare, pharma, banking, utilities, and public sector organizations. We help them launch and secure AI initiatives while protecting crown-jewel data, meeting regulatory expectations, and keeping cyber risk under control.



Representative clients (anonymized)

MONIXITY
Securing Today. Ensuring Tomorrow.

Next Step:
Secure Your AI Initiatives

If this checklist surfaced gaps or opportunities in your AI security program, let's map out your next moves. Monixity helps CISOs and technology leaders secure AI systems, protect crown-jewel data, and move fast without losing control.

- **Review your AI Security Readiness score**
- **Identify your top 2–3 priority actions**
- **Discuss a roadmap for secure AI pilots and scale-out**

Contact Us

www.monixity.com
info@monixity.com
LinkedIn: linkedin.com/company/monixity
X: x.com/monixity