Arx Certa
Expert UK Cloud, Infrastructure & AI

AI READINESS · SECURITY

Secure AI Adoption Checklist

30 security-anchored items every UK firm should be able to evidence before AI use becomes production-grade.

Free download · May 2026 · v2026-05
Scorecard companion: **arxcerta.com/ai-readiness-scorecard**

Why a security-first checklist

Every AI adoption discussion is also a security discussion, whether the people in the room know it yet or not. New data flows, new vendors, new decision-makers (the AI itself, with human review on top), and new categories of risk surface. This checklist is the security baseline an AI programme should meet — drawn from NCSC guidance, ISO 27001 controls, and the patterns Arx Certa sees in real UK engagements.

> **How to use.** Six items across five security dimensions. Tick where evidence exists today. The gaps are the work that needs to happen before AI moves from pilot to production.

Identity and access

- ☐ MFA is enforced on every AI tool account.
- ☐ SSO routes all AI tool authentication through the corporate IdP.
- ☐ Role-based access control restricts AI tools to the cohorts that need them.
- ☐ Privileged accounts are excluded from broad AI tool access until reviewed.
- ☐ Offboarding immediately revokes AI tool access at the IdP level.
- ☐ Service accounts used by AI tools have least-privilege scopes.

Data protection

- ☐ Sensitivity classification has been applied to AI-touched data.
- ☐ Encryption at rest and in transit is verified for AI tool integrations.
- ☐ DLP policies cover AI inputs and outputs, not just classic file flows.
- ☐ Data residency is contractually confirmed where regulated.
- ☐ Retention policies apply consistently to AI inputs and outputs.
- ☐ Backup and restore covers AI-generated content.

Network and infrastructure

- ☐ Network segmentation isolates production AI tooling from general user network.
- ☐ API integrations between AI tools and core systems pass through reviewed gateways.
- ☐ Logging captures AI tool API activity (who, when, what).
- ☐ Rate limiting protects against runaway agent loops.
- ☐ DR plan covers systems AI use cases depend on.
- ☐ Capacity is sufficient for peak AI workloads without degrading other services.

Vendor and supply chain

- ☐ Each AI vendor is mapped in the third-party register.
- ☐ DPAs are signed with every AI vendor processing personal data.
- ☐ SOC2 / ISO27001 evidence has been reviewed for every Tier 2+ AI vendor.
- ☐ Sub-processor disclosures are current and acknowledged.
- ☐ Right-to-audit clauses are present in vendor agreements where material.
- ☐ Vendor incident notification times match company breach-reporting obligations.

Monitoring, IR, and audit

- ☐ AI tool audit logs are retained for at least 12 months.
- ☐ Incident response plan covers AI-specific scenarios (prompt leak, jailbreak, output bias).

- ☐ Anomaly detection covers AI tool usage (unusual volumes, unusual prompts).
- ☐ Pen testing scope explicitly includes AI integrations.
- ☐ Annual security review covers the AI tooling estate.
- ☐ Audit-ready evidence pack can be produced within 7 days.

What to do with this

Tick what you can. For everything else, the gaps are the security work that should happen before AI use becomes production-grade. The Arx Certa AI Readiness Scorecard scores Security as one of five dimensions.

Want a score, not just a checklist?

The 4-minute scorecard turns this checklist into a weighted score across 5 dimensions, a readiness band, and a personalised 30-day action plan you can take to your leadership team.

Take the scorecard → arxcerta.com/ai-readiness-scorecard