ISO/IEC 27090:2026-INFORMED

AI is in your software.
Is security in your rollout?

A practical AI adoption security checklist

For organisations adding AI to existing applications or using software with AI features.

20 CHECKS
Evidence you can inspect

→

5 STAGES
From approval to operation

→

1 USE CASE
Assess each feature separately

USE IT WITH YOUR APP OWNER + SECURITY TEAM

Tick only when evidence supports the answer. For every unticked item, record a gap, an unassessed status, or a reason it does not apply. Track an owner and due date.

For hosted software, request product-specific evidence from the supplier when you cannot test a control yourself.

Scope: AI users and application integrators. Excludes building or training models. Selected security checks, not an exhaustive audit or certification checklist. Public-source basis and limits: page 7.

01 / APPROVE

Know what you are enabling.

Start before procurement, a pilot, or switching on an AI feature.

BUSINESS
Purpose + impact

→

SUPPLIER
Evidence + limits

→

OWNER
Accepts the risk

☐ 01 **Name the use case** G | I §6

Does the register identify this AI feature, its purpose and accountable owner?

Evidence: Feature register and named business owner.

☐ 02 **Set the boundary** G | I §6

Is the approved use explicit about the decisions this AI may influence?

Evidence: Use-case approval; excluded decisions and consequences.

☐ 03 **Review the supplier** G | A: secure-by-design

Has the supplier provided security evidence for the AI feature you will use?

Evidence: Product-specific testing, limitations and dependencies.

☐ 04 **Allocate responsibility** G | A: something goes wrong

Are customer and supplier incident-response duties agreed?

Evidence: Contract or responsibility record with escalation contacts.

I §6 = practical first steps. A = adopter guidance; source key on page 7.

02 / PROTECT DATA

Control what AI can see.

Prompts, uploaded files and connected records all need boundaries.

INPUT
Approved information

→

RETRIEVAL
User-level access

→

OUTPUT
No unintended disclosure

☐ **05 Approve data use** G | A: privacy / data

Is supplier reuse of your inputs for model training explicitly accepted or disabled?

Evidence: Current terms plus the tenant setting.

☐ **06 Limit what goes in** G | O02: mitigation

Do controls prevent unapproved sensitive data from entering AI requests?

Evidence: Test masking or blocking with synthetic sensitive records.

☐ **07 Preserve permissions** G | O08: mitigation 1

If AI searches company files, can it retrieve only what that user may access?

Evidence: Cross-user access test, including search results and citations.

☐ **08 Control knowledge sources** G | O08: mitigation 2

If AI reads a knowledge base, are source changes checked before ingestion?

Evidence: Approved sources and a rejected tampered-document test.

A: privacy / data. O02 and O08: OWASP mitigations; source key on page 7.

03 / LIMIT AUTHORITY

Give AI only the access it needs.

An AI answer should never grant itself permission to act.

IDENTITY
Authenticate the user

→

PERMISSION
Enforce in the app

→

APPROVAL
Before high-impact action

☐ 09 **Protect sign-in** G | A: authentication

Is multi-factor authentication enforced for access to the AI service?

Evidence: Sign-in policy and an observed login test.

☐ 10 **Constrain connectors** G | O06: mitigation 1-5

If AI can call tools, are permissions limited to the approved task?

Evidence: Connector scopes; remove unused tools and broad write access.

☐ 11 **Check every action** G | O06: mitigation 7

If AI calls another system, does that system enforce the user's permissions?

Evidence: A forbidden action is denied even when AI requests it.

☐ 12 **Gate high-impact actions** G | O06: mitigation 6

If AI can send, delete or change records, do high-impact actions require approval?

Evidence: A workflow test showing approval before execution.

A: authentication. O06: OWASP Excessive Agency; source key on page 7.

04 / TEST

Try to break the boundaries.

Use a safe test environment and synthetic data. Record failures.

ATTACK
Malicious input

→

BOUNDARY
Data + actions

→

EVIDENCE
Pass, fail, remediate

☐ **13 Challenge malicious prompts** — G | O01: mitigation 7

For generative AI, have direct and document-borne prompt attacks been tested?

Evidence: Safe tests show whether data or action boundaries fail.

☐ **14 Validate downstream output** — G | O05: mitigation

If AI output enters code or another app, is it checked before processing?

Evidence: Tests for unsafe commands, queries and rendered content.

☐ **15 Test manipulated inputs** — G+P | I §4 + P

For predictive or classification AI, have realistic altered inputs been tested?

Evidence: Security test results and recorded failure limits.

☐ **16 Check important answers** — G | A: staff / output validation

Are consequential AI outputs verified before people act on them?

Evidence: A sampled decision with source checks or expert review.

O01 / O05: OWASP. I §4: threats. A: output validation. P: added practice.

05 / OPERATE

Keep control after go-live.

AI features and supplier behaviour change. Revisit the evidence.

DETECT Observe + investigate → **CONTAIN** Disable + fall back → **REVIEW** Retest + improve

☐ **17 Detect suspicious activity** G+P | I §4 + P

Can monitoring reveal unusual queries, data access or behaviour changes?

Evidence: Alert test and a named responder; protect sensitive logs.

☐ **18 Recheck after change** G+P | I §4 + P

Are safeguards retested when the provider, model or integration materially changes?

Evidence: Change review and repeated security tests.

☐ **19 Practise stopping safely** P | P: response drill

Has the team tested disabling the AI feature and using a fallback?

Evidence: Exercise record: stop access, preserve evidence, continue work.

☐ **20 Plan the exit** P | P: exit control

Can the organisation remove its data and revoke integrations when use ends?

Evidence: Deletion terms and a tested account / connector offboarding path.

I §4: protection over time. P: MicroApz operational practices; page 7.

06 / DECIDE

Turn checks into a decision.

A tick count is not a security rating. Review the consequences of each gap.

HOLD

Uncontrolled exposure or action: resolve before use.

RESTRICT

Reduce data or permissions while closing gaps.

PROCEED

Owner accepts residual risk; monitoring is ready.

RECORD THE DECISION / MICROAPZ ADDED PRACTICE

AI feature / use case

Gap IDs + evidence location

Action owner + due date

Decision + approver + review date

SOURCE KEY & COVERAGE

ISO/IEC 27090:2026 provides AI security guidance. This independent aid uses ISO public summaries, not the full standard. No ISO clause conformity is claimed. Model development, training and detailed model-theft defences are outside this adopter shortlist.

I	ISO public explainer, §§4 & 6; based on the final draft.
2026	Published edition confirmed in the ISO catalogue (Oct 2026).
A	Engaging with artificial intelligence (24 Jan 2024), mitigations.
Oxx	OWASP LLM Top 10 (2025), numbered risk / mitigation sections.
P	MicroApz added practice; not an ISO requirement.

G = guidance; P = added practice. Small labels on each check identify its basis and source section. Keep existing cybersecurity controls in place. Links work after download. Checked 2 November 2026.