SAVVIES DESK
INTELLIGENCE FOR DECISION MAKERS

NEW STANDARD · OCTOBER 2026

ISO/IEC 27090

A new ISO standard for AI security.

What does it change for your organisation?

WHAT IT IS

A threat guide for AI systems. Not a certificate.

It explains how attacks specific to AI work, what they do to the business, and how to detect and reduce them across the whole AI life cycle.

PUBLISHED
14 Oct 2026

SIZE
Edition 1 · 56 pages

AUTHOR
JTC 1/SC 27, the 27001 committee

TYPE
Guidance. Nobody certifies to it.

Its job is protecting AI systems, not using AI to defend networks.

WHERE IT FITS

42001 proves you govern AI. 27090 shows how it gets attacked.

27001 — **Is our ISMS working?**
Requirements · certifiable

27002 — **Which security controls apply?**
Control guidance

42001 — **Is AI governed responsibly?**
AI management system · certifiable

27090 — **Which AI threats exist, and how do we counter them?**
Threat and mitigation guidance

42001 Annex A asks for security measures but names no attacks. 27090 is the catalogue it was missing, much as 27017 is for cloud.

HOW IT IS BUILT

Three clauses do the work.

1 **Applying security to AI**

Governance, zero trust, secure engineering, supply chain, threat modelling and red teaming.

2 **Types of compromise**

Each threat gets a description, its consequences, how to detect it and how to mitigate it.

3 **Mitigations**

Including how they interact, conflict with each other and lose effect over time.

ANNEXES

A: attacks mapped to life cycle and assets · B: AI versions of conventional attacks · C: federated learning

CLAUSE 6 · THE THREATS

Nine ways an AI system gets compromised.

6.2

Data poisoning

Training data altered to change behaviour

6.3

Evasion

Crafted inputs make the model misbehave

6.4

Membership inference

Learning whether a record was in the training data

6.5

Model exfiltration

Rebuilding a model by querying it

6.6

Model inversion

Reconstructing training data from outputs

6.7

Direct model poisoning

The model itself tampered with in build or operation

6.8

Sensitive output

Responses leak prompts, inputs or training data

6.9

Prompt injection

Instructions smuggled in, even via emails and documents

6.10

Output injection

Model output carries a payload a downstream system runs

Most of these never touch a firewall rule or a vulnerability scan.

SAVVIES DESK
INTELLIGENCE FOR DECISION MAKERS

KEY CONTROLS

Controls follow the AI life cycle.

01	**Data**	**Provenance and lineage records** Counters data poisoning
02	**Training**	**Secured build environment and model supply chain** Counters direct model poisoning
03	**Evaluation**	**Red teaming and adversarial testing** Counters evasion and inference attacks
04	**Deployment**	**Least privilege, input and output handling** Counters prompt and output injection
05	**Monitoring**	**Logging, rate limits, query pattern alerts** Counters model exfiltration and abuse
06	**Retirement**	**Inventory closure and secure disposal** Counters leftover weights, data and logs

For every threat: a named owner, an enforcement point, and evidence the control actually works.

WHAT CHANGES IN YOUR THINKING

Three shifts security teams should make.

01 **Today's input is tomorrow's training data.**

If production data feeds a retrain, poisoning can arrive long after development.

02 **A secure platform can still give a compromised answer.**

A poisoned ticket or document arrives through a legitimate channel. No breach, no alert.

03 **Every approval has an expiry condition.**

A new model version, data source or tool permission means the old test no longer counts.

KNOW THE LIMITS

Where the standard stops.

No certificate, no shortcut

Use it inside a 27001 or 42001 scope. It is not a harmonised standard, so it gives no EU AI Act presumption of conformity on its own.

Four years from proposal to print

Drafted in the image classifier era, so agents, tool calling and MCP get thin treatment, and jailbreaks have no section of their own.

Some fixes are still research

Several mitigations are research results rather than proven operational practice. Accidental and insider failures get less attention.

PAIR IT WITH

OWASP Top 10 for LLM Applications and MITRE ATLAS

BANKS · SACCOS · INSURERS

Your AI risk probably sits inside a vendor contract.

Chatbots, credit scoring, KYC and fraud engines arrive as features of platforms you already approved. The platform's certificate says nothing about how the AI behaves under attack.

FOUR QUESTIONS FOR EVERY AI VENDOR

Q1 Is prompt injection in scope of your vulnerability disclosure?

Q2 How are model weights and training data protected?

Q3 Which inference and tool call logs will we receive?

Q4 Who approves model updates, and do we get notice?

Membership inference and inversion are confidentiality failures. Where personal data trained the model, data protection law applies.

START HERE THIS QUARTER

Four moves. No new programme needed.

1

Inventory every AI system, including AI inside vendor tools.

2

Add 27090 to your Statement of Applicability as a control source, as 27017 is for cloud.

3

Threat model and red team your highest impact AI system.

4

Keep evidence as part of operations, not audit preparation.

Save this for your next risk committee. Which AI system would you threat model first?

From the Desk