This is a preview of ISO/IEC 27090:2026. Click here to purchase the full version from the ANSI store.

International

ISO/IEC 27090

First edition
2026-10

Cybersecurity — Artificial intelligence (AI) — Addressing security threats and compromises to AI systems

Cybersécurité — Intelligence artificielle (IA) — Traitement des menaces et des situations susceptibles de compromettre les systèmes d'IA

Reference number
ISO/IEC 27090:2026(en)

© ISO/IEC 2026

This is a preview of ISO/IEC 27090:2026. Click here to purchase the full version from the ANSI store.



COPYRIGHT PROTECTED DOCUMENT

© ISO/IEC 2026

All rights reserved.

This ISO publication is protected by copyright and is owned by ISO and/or its licensors.

The content of this ISO publication is provided under licence, not sold. Use is subject to the applicable licence terms issued by ISO, an ISO member body, or an authorized third-party distributor.

Except as required for implementation or expressly permitted by a separate licence, no part of this ISO publication may be reproduced, distributed, modified, used, or made available in any form or by any means – electronic or mechanical, including photocopying, scanning, recording, or posting on internal or external digital platforms.

Any use beyond the scope of the granted rights is prohibited and may result in legal action.

ISO copyright office
CP 401 • CH-1214 Vernier, Geneva
Phone: +41 22 749 01 11
Email: copyright@iso.org
Website: www.iso.org

Published in Switzerland

This is a preview of ISO/IEC 27090:2026. Click here to purchase the full version from the ANSI store.

Foreword ..... **v**

Introduction ..... **vi**

1 Scope **1**

2 Normative references **1**

3 Terms and definitions **1**

4 Abbreviated terms **5**

5 Application of information security **5**
5.1 General 5
5.2 AI governance programme 6
5.3 Zero trust principles 7
5.4 Applying software engineering best practices to AI 8
5.5 Supply chain management 8
5.5.1 General 8
5.5.2 AI-specific supply chain threats 9
5.5.3 Complex AI supply chains 9
5.5.4 Specific attack vectors 9
5.5.5 Mitigation strategies 9
5.6 Threat modelling 10
5.7 Red teaming 10

6 Types of potential compromise to AI systems **10**
6.1 General 10
6.2 Data poisoning attack 13
6.2.1 Description 13
6.2.2 Failures resulting from data poisoning attacks 14
6.2.3 Detection of data poisoning attacks 14
6.2.4 Mitigation of data poisoning attacks 15
6.3 Evasion attack 17
6.3.1 Description 17
6.3.2 Failures resulting from evasion attacks 18
6.3.3 Detection of evasion attacks 18
6.3.4 Mitigation of evasion attacks 18
6.4 Membership inference 20
6.4.1 Description 20
6.4.2 Failures resulting from membership inference 20
6.4.3 Detection of membership inference 21
6.4.4 Mitigation of membership inference 21
6.5 Model exfiltration 21
6.5.1 Description 21
6.5.2 Failures resulting from model exfiltration 22
6.5.3 Detection of model exfiltration 22
6.5.4 Mitigation of model exfiltration 23
6.6 Model inversion 23
6.6.1 Description 23
6.6.2 Failures resulting from model inversion 23
6.6.3 Detection of model inversion 24
6.6.4 Mitigation of model inversion 24
6.7 Direct model poisoning 24
6.7.1 Description 24
6.7.2 Failures resulting from direct model poisoning 25
6.7.3 Detection of direct model poisoning 25
6.7.4 Mitigation of direct model poisoning 25
6.8 Sensitive model output 25
6.8.1 Description 25

© ISO/IEC 2026 – All rights reserved

This is a preview of ISO/IEC 27090:2026. Click here to purchase the full version from the ANSI store.

6.8.4 Mitigations of sensitive model output ... 25
6.9 Prompt injection ... 26
6.9.1 Description ... 26
6.9.2 Failures resulting from prompt injection ... 26
6.9.3 Detection of prompt injection ... 26
6.9.4 Mitigation of prompt injection ... 27
6.10 Output contains injection attacks ... 27
6.10.1 Description ... 27
6.10.2 Failures resulting from output contains injection attacks ... 27
6.10.3 Detection of output contains injection attacks ... 27
6.10.4 Mitigations of output contains injection attacks ... 27
6.11 Considerations for AI systems ... 28
6.11.1 General ... 28
6.11.2 Reduce impact by data limitation ... 28
6.11.3 Governing and controlling the effect of unwanted model behaviour during the operation stage ... 29
6.11.4 AI system and data provenance ... 32

7 Mitigations and their interactions with threats and other applied mitigations ... 33
7.1 General ... 33
7.2 Conflicting interactions of mitigations ... 33
7.3 Continuity of mitigations across the AI life cycle ... 33
7.4 Model and mitigation deterioration over time ... 33
7.5 Logging and monitoring ... 34
7.6 Development environment protection ... 34
7.7 Malicious input detection and filtering ... 35
7.7.1 Detecting out-of-distribution inputs ... 35
7.7.2 Detecting adversarial inputs ... 35
7.8 Throttling model use ... 36
7.9 Optimising the size of the model ... 36

Annex A (informative) **Mapping attacks to AI system life cycle and to assets ... 37**

Annex B (informative) **AI-specific versions of conventional attacks ... 44**

Annex C (informative) **Federated machine learning ... 49**

Bibliography ... 51

© ISO/IEC 2026 – All rights reserved

This is a preview of ISO/IEC 27090:2026. Click here to purchase the full version from the ANSI store.

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular, the different approval criteria needed for the different types of document should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see www.iso.org/directives or www.iec.ch/members_experts/refdocs).

ISO and IEC draw attention to the possibility that the implementation of this document may involve the use of (a) patent(s). ISO and IEC take no position concerning the evidence, validity or applicability of any claimed patent rights in respect thereof. As of the date of publication of this document, ISO and IEC had not received notice of (a) patent(s) which may be required to implement this document. However, implementers are cautioned that this may not represent the latest information, which may be obtained from the patent database available at www.iso.org/patents and https://patents.iec.ch. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation of the voluntary nature of standards, the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the World Trade Organization (WTO) principles in the Technical Barriers to Trade (TBT) see www.iso.org/iso/foreword.html. In the IEC, see www.iec.ch/understanding-standards.

This document was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 27, *Information security, cybersecurity and privacy protection*.

Any feedback or questions on this document should be directed to the user's national standards body. A complete listing of these bodies can be found at www.iso.org/members.html and www.iec.ch/national-committees.

© ISO/IEC 2026 – All rights reserved

This is a preview of ISO/IEC 27090:2026. Click here to purchase the full version from the ANSI store.

Mitigating the risks of security threats to AI systems in a timely manner can help to improve the trustworthiness of AI systems and of the organizations that use or develop them, and thus earn trust in these systems and organizations. As technology continues to advance and innovate at a fast pace, emerging and new security vulnerabilities discovered in AI systems can require continuous evaluation of security objectives.

AI systems are subject to all of the traditional cybersecurity vulnerabilities, exploits and attacks. They can be subject to new threats, notably those based on data, input manipulation or actions of AI models. However, these additional concerns do not supersede conventional cybersecurity concerns. As the overall cybersecurity threat landscape evolves, the need to protect AI systems against non-AI system specific threats continues to be an ongoing pursuit.

AI systems have become more widespread as organizations embrace digital transformation, resulting in an increasing likelihood of security attacks on these systems. Furthermore, as a result of the disparate and widespread use of AI systems, and especially within safety critical contexts, the result of security attacks can be severe and in some cases, lead to physical or mental repercussions for individuals.

Due to the way AI systems are developed, and their strong dependency on training data, which can include user data or unvetted data, these systems can have additional vulnerabilities compared to conventional information processing systems. Furthermore, there is a lack of understanding when it comes to the inner workings of AI, specifically when it comes to understanding and diagnosing individual instances of AI behaviour. This document provides guidance to organizations who are choosing to develop or deploy AI systems, providing awareness of the types of security threats such systems can face, and helping organizations understand the consequences of such threats and the most commonly related detection and mitigation techniques. Such awareness and understanding of this new threat landscape is imperative for securing AI systems against adversarial attacks.

This document describes various security threats and the exploitation of vulnerabilities that can happen when using AI systems. The objective is to identify recommended measures to improve the security of AI systems by identifying threats and the consequences of these threats. The identification of threats can help users and developers of AI systems better understand security issues related to AI systems.

Security threats and vulnerabilities with a lack of controls in AI systems can lead to failures, impacting both individuals and businesses.

It can be challenging to detect some threats to AI systems and thus, developing controls without effective threat detection would not be useful or adequate. This document can help identify where security threats can be effectively addressed.

In general, it is recommended that the established commonly accepted security practices in ISO/IEC 27001[1] and ISO/IEC 27002[2] are adhered to. However, the reliance of many AI systems on training data introduces new threats. Attacks such as data poisoning can cause undesirable outcomes by injecting malicious or unwanted data into training data. During the operational phase, AI systems gather data to support future model updates. However, this process can be vulnerable to the introduction of corrupted data, potentially compromising the integrity of future model retraining during the design and development stage. Where continuous or iterative learning occurs, data inputs supplied during operation can also become training data, so corrupted inputs can also become a vector for data poisoning. Accordingly, AI or machine-learning-enabled work, data, and control flows can benefit from some additional controls. This document can provide the necessary insights to apply both established and new controls, for example by explaining what the assets are, and what aspects need protecting.

Other attacks such as model exfiltration through querying can be mitigated by conventional information security measures such as rate limiting. Models and intellectual property can be stolen even when effective conventional information security measures are in place, since a model exfiltration attack is established through legitimate querying of the model. Threat detection is a challenge but an essential component to effectively mitigate model exfiltration. The measures used to mitigate AI threats can use a defence-in-depth strategy that leverages multiple security measures to protect an organization's assets.

© ISO/IEC 2026 – All rights reserved

This is a preview of ISO/IEC 27090:2026. Click here to purchase the full version from the ANSI store.

© ISO/IEC 2026 – All rights reserved