

HUMAN OVERSIGHT · ACCOUNTABILITY · RISK PROPORTIONALITY

AI supports people. Accountability stays human.

Principles for the responsible development and use of artificial intelligence.

DOCUMENT ID	VERSION	CLASSIFICATION	DOCUMENT OWNER	REVIEW CYCLE	PUBLISHED
ITO-TRUST-AI-001	1.0	PUBLIC	Engineering & Information Security	Annual	October 2026

Introducing AI responsibly

Artificial intelligence can significantly improve how organisations process information, identify patterns, automate repetitive activities and support operational decision-making.

In government, public safety and other mission-critical environments, these capabilities must be introduced responsibly. The potential benefits of AI must be considered together with questions of security, privacy, accountability, transparency, reliability and human oversight.

ITO approaches artificial intelligence as an **engineering capability** that supports defined operational objectives within appropriate governance frameworks.

This document establishes the principles that guide ITO when designing, developing, integrating or deploying AI-assisted functionality.

CONSIDERED TOGETHER WITH THE BENEFITS OF AI

- Security
- Privacy
- Accountability
- Transparency
- Reliability
- Human oversight

CONTENTS

01	Purpose	02
02	Our approach to AI	03
03	Responsible AI principles	04
04	Mission-critical environments	05
05	Human-in-the-loop	06
06	AI is not evidence by default	07
07	AI lifecycle governance	08
08	Third-party AI	08
09	Responsible AI & regulation	09
10	Our commitment	09

DOCUMENT CONTROL


DOCUMENT ID	ITO-TRUST-AI-001
VERSION	1.0
CLASSIFICATION	PUBLIC
OWNER	Engineering & Information Security
REVIEW CYCLE	Annual
PUBLISHED	October 2026

AI should support people, not remove accountability

ITO develops technology to support people and institutions in performing their responsibilities more effectively.

Where AI-assisted functionality is used in workflows involving consequential decisions, the appropriate level of human oversight should be determined according to the nature, risk and legal context of the use case.

TECHNOLOGY MAY ASSIST WITH



AUTOMATION AND CONTROL

The question is not whether a task is automated.

It is whether decisions with serious consequences remain subject to **appropriate control**.

Six principles

01

Human oversight

AI-assisted systems should preserve an appropriate level of human oversight according to the nature and potential impact of the use case. For higher-impact workflows, technology should support authorised personnel rather than obscure or eliminate institutional responsibility.

AI can support judgement. Accountability remains human and institutional.

02

Accountability

Every AI capability should exist within a defined operational and governance context. Institutions remain responsible for determining how technology is authorised and used within their processes, while ITO is responsible for engineering and maintaining technology within its agreed scope.

Responsibility should remain identifiable even when individual activities are automated.

03

Transparency

Users should have an appropriate understanding of when AI-assisted functionality contributes to an operational process. Where technically feasible and relevant to the use case, systems should provide information that supports understanding of AI-assisted outputs and their role within the broader workflow.

04

Privacy & data governance

AI does not create a new justification for processing information. The use of personal or sensitive information within AI-assisted functionality must remain subject to applicable legal authority, purpose limitation, access controls and data-governance requirements.

Technical capability does not create legal authority.

05

Security & reliability

AI functionality forms part of the overall security architecture of the system in which it operates. Relevant considerations may include protection of input and output data, access control, model and service integrity, system availability, dependency security and monitoring.

AI-assisted functionality should be evaluated according to the operational consequences of incorrect, unavailable or manipulated outputs.

06

Risk proportionality

Not all AI applications create the same level of risk. An AI capability used to categorise technical logs does not present the same considerations as technology used within an operational public-safety workflow.

Governance, testing, documentation and human oversight are proportionate to the nature and potential impact of the AI use case.

Questions before introducing AI

The use of AI in government and public safety requires additional consideration because technology may operate within processes affecting individuals, institutional decisions or operational activities.

BEFORE INTRODUCING AI-ASSISTED FUNCTIONALITY, RELEVANT CONSIDERATIONS MAY INCLUDE

Purpose?

What problem is the AI capability intended to solve?

Authority?

Is the intended processing permitted within the applicable institutional and legal framework?

Data?

What information is required and how is it governed?

Impact?

What could happen if the AI output is incorrect?

Oversight?

What level of human review is appropriate?

Traceability?

Can relevant system activity be reconstructed where required?

Security?

How is the AI component protected and monitored?

Human oversight by design

Human-in-the-loop mechanisms can be an important safeguard where AI supports workflows with meaningful operational or individual consequences.

DEPENDING ON THE USE CASE, AN
AUTHORISED USER MAY

01 Review an AI-generated result

02 Validate identified information

03 Confirm a proposed action

04 Reject or correct an output

05 Make the final operational
decision

The appropriate model depends on the specific
functionality and risk.

MEANINGFUL REVIEW


**ITO does not treat
human oversight
as a cosmetic
confirmation step.**

Where human review is required,
the workflow should provide the
user with sufficient information
and authority to meaningfully
assess the AI-assisted output.

Distinguishing detection from determination

AI-generated detections, classifications, confidence scores or recommendations should not automatically be treated as independently verified facts.


Their evidentiary or operational significance depends on the technology, workflow, validation process and applicable legal framework.



**Detection is not
the same as
determination.**

Responsibility continues after deployment

Responsible AI does not end when a model or AI service is deployed. ITO's approach considers the lifecycle of AI-assisted functionality:




Depending on risk and technical architecture, lifecycle activities may include documenting the intended purpose, assessing data requirements, testing relevant performance characteristics, evaluating failure scenarios, controlling deployment changes and reviewing operational behaviour.

External models and AI services

ITO solutions may incorporate or integrate AI technologies provided by third parties. **The use of an external AI technology does not remove ITO's responsibility to assess its suitability for the intended integration within ITO's scope.**

RELEVANT CONSIDERATIONS MAY INCLUDE



Particularly sensitive institutional information should not be transmitted to external AI services without appropriate **technical, contractual and governance safeguards**.

An evolving regulatory landscape

ITO monitors the evolution of AI governance, technical standards and regulatory requirements relevant to the environments in which its technology is used. Our responsible AI approach may reference recognised frameworks and principles including:

REFERENCED

EU AI Act

Where applicable

NIST AI RMF

AI Risk Management Framework

OECD AI Principles

International principles

ISO/IEC 42001

And related AI management standards



References to external frameworks do not constitute certification or a universal statement of compliance. Applicability depends on the technology, use case, jurisdiction and deployment context.

10 • OUR COMMITMENT

Artificial intelligence creates significant opportunities to improve digital government and mission-critical operations. Those opportunities must be pursued responsibly. ITO is committed to developing AI-assisted technology with appropriate consideration for human oversight, accountability, privacy, security, transparency and risk.

NOT OUR OBJECTIVE

To automate every decision that technology makes possible.

OUR OBJECTIVE

To use technology where it creates **meaningful value** while preserving the **governance and accountability** required by the environment in which it operates.

**Responsible AI is not about limiting innovation.
It is about making **innovation trustworthy**.**



ITO TRUST DOCUMENTS

Engineering Digital Trust.

RELATED DOCUMENTS

ITO-TRUST-SEC-001

ITO Security Whitepaper

ITO-TRUST-PRV-001

ITO Privacy & Data Governance

TRUST DOCUMENT

ITO Secure Development Overview

DOCUMENT

ITO Responsible AI Principles

ID · VERSION · CLASSIFICATION

ITO-TRUST-AI-001 · 1.0 · PUBLIC

OWNER · REVIEW

Engineering & Information Security · Annual · October 2026

CONTACT

security@ito.dev · ito.dev/en/trust-center

COMPANY

IT Odjel Mostar d.o.o., Mostar, Bosnia and Herzegovina