
Zorlu Holding · Technology and Digital Business Development Group Presidency

AI GOVERNANCE FRAMEWORK

Public Summary

CONTENTS

Contents

01	Why an AI Governance Framework?	3
02	Our Core Approach and Reference Standards	4
03	Core Rules for Every Employee	5
04	Risk Classification Matrix	6
05	Responsible and Ethical AI Principles	7
06	AI Governance Structure	8
07	AI Project Lifecycle	10
08	Data Security and Control Approach	11
09	Agent and MCP Governance	12
10	Employee Awareness and Shadow AI Prevention	13
11	Incident Response & Business Continuity	14
12	Audit, Targets and Continuous Improvement	15
	Glossary	16

01 · WHY AN AI GOVERNANCE FRAMEWORK?

Why an AI Governance Framework?

For most organisations, the real question is no longer whether to use AI, but how to use it safely, ethically and lawfully.

AI is rapidly transforming the way organisations operate. While this transformation creates significant opportunities, it also introduces new risks in areas such as data security, personal data protection, discrimination, intellectual property and reputation.

At Zorlu Holding, we have established an AI Governance Framework to answer this question systematically. The framework defines the principles that govern AI systems across their entire lifecycle, from ideation to decommissioning; who is accountable for what; how risks are classified and managed; and how reporting works.

The following sections summarise the key components of the framework: risk classification aligned with international standards, responsible AI principles, governance structure, project lifecycle, data security, agent and MCP governance, employee awareness and the prevention of shadow AI.

ABOUT THIS DOCUMENT

This document is the public summary of Zorlu Holding's internal AI Governance Framework. It outlines our approach to corporate AI governance and aims to offer a practical reference for fellow professionals in the field. Organisation-specific implementation parameters, internal control thresholds and operational details have been excluded from this version. This document does not constitute legal advice.

POSITIONING NOTE

Some practices in this framework are already in place across the organisation, while others are defined as target states. We follow the framework as a guide for raising our maturity step by step. By sharing this document publicly, we aim to offer a useful reference to organisations and peers on the same journey.

02 · OUR CORE APPROACH AND REFERENCE STANDARDS

Our Core Approach and Reference Standards

The framework draws not on a single source but on internationally recognised frameworks. This allows us both to align with global best practice and to stay prepared for regulatory developments.

EU AI Act

The unacceptable, high, limited and minimal risk tiers map one-to-one to our framework's risk categories. For high-risk systems, we reference the requirements of Articles 9–15.

OECD and UNESCO

The foundation of our responsible AI principles; human oversight, transparency, fairness and accountability are aligned with them.

KVKK · GDPR · SMK · FSEK

Data protection under KVKK (Law No. 6698) and the GDPR, and intellectual property obligations under SMK (Law No. 6769) and FSEK (Law No. 5846), are an integral part of the process.

OWASP LLM Top 10 and MCP security

These form the basis of our security training and controls for developers and data scientists.

CRISP-DM, TDSP, OSEMN

We build our project and data science processes on these established methodologies; our project lifecycle is aligned with them.

Türkiye AI Action Plan (2026–2030)

The axes and principles of the plan, enacted by the Presidential Circular of 18 August 2026 (ethical responsibility, digital sovereignty, sustainability, widespread adoption), are aligned with our framework.

THE ESSENCE OF OUR APPROACH

Managing risk across the lifecycle: from the moment an AI idea emerges, its risk is assessed, controls proportionate to the risk level are applied, and the system continues to be monitored after it goes live.

03 · CORE RULES FOR EVERY EMPLOYEE

Core Rules for Every Employee

On the ground, a good governance framework translates into simple rules that every employee can keep in mind in their daily work.

We designed the core rules below to be memorable even without reading the full document.

TOPIC	CORE RULE
Approved tool use	Only AI tools approved by the organisation are used, and only with corporate accounts. Company data is never processed through personal accounts.
Confidential data input	Confidential or restricted data is not entered into publicly available AI tools. This includes customer, financial and employee data as well as trade secrets.
Personal data	Personal data such as ID numbers, IBANs, phone numbers, e-mail addresses and health data is masked or anonymised before being entered into any system.
AI in Decision-Making	Decisions on hiring, promotion, performance, financial approval or health are never left to AI. Output serves as decision support; the final decision rests with an authorised person.
Output accuracy	AI output is verified against its source before it is published or used as the basis for a decision. AI can hallucinate.
Transparency	AI-generated content (images, text, code) is clearly labelled as such.
Incident reporting	Any suspected data leak, erroneous or discriminatory output, or unexpected behaviour is logged through the incident reporting channel.

04 · RISK CLASSIFICATION MATRIX

Risk Classification Matrix

The framework classifies all AI use cases into four categories that map one-to-one to the EU AI Act's four-tier approach. Every project is assessed against these categories at the ideation stage using a Risk Assessment Form; the higher the risk, the more intensive the controls and the shorter the deadline for closing findings.

CATEGORY	DEFINITION AND EXAMPLES	MANDATORY CONTROLS
A · Unacceptable	Prohibited under AI Act Art. 5: social scoring, subliminal manipulation, workplace emotion recognition, personality or crime-propensity profiling.	Fully prohibited. Any development, procurement or use is stopped immediately.
B · High Risk	HR decisions, critical infrastructure, access decisions (credit, insurance, health, education), biometric categorisation.	FRIA and DPIA, full technical documentation, human oversight, log retention, OWASP LLM Top 10 controls.
C · Limited Risk	Systems that interact directly with users, content-generating systems, decision support systems.	AI Act Article 50 transparency obligation; users are informed and synthetic content is labelled.
D · Low Risk	Operational support, code completion, internal drafting, standard analytics.	Standard security and data classification rules; inventory registration.

WHY IT MATTERS

Mapping risk categories one-to-one to an international regulation ensures internal consistency and reduces the compliance burden as regulation evolves. This mapping makes the framework both auditable and ready for future regulation.

05 · RESPONSIBLE AND ETHICAL AI PRINCIPLES

Responsible and Ethical AI Principles

We have defined eight core principles to ensure that AI systems operate fairly, explainably and under human oversight. These principles apply at every stage of the lifecycle and are aligned with the OECD and UNESCO frameworks.

01 Human oversight and control

In human rights, employment, financial approval and health decisions, AI output is decision support only; an authorised person makes the final decision and records the rationale.

03 Fairness and non-discrimination

No discrimination on grounds such as gender, age, religion, language, race or disability; models are regularly tested for performance gaps across subgroups.

05 Accountability

Every AI system has a designated business owner, risk owner and technical owner.

07 Sustainability

The energy consumption of large models is monitored; smaller, more efficient models delivering the same business value are preferred; carbon impact is reported annually.

02 Transparency and explainability

Users know when they are interacting with an AI system (AI Act Art. 50); synthetic content is labelled; a right to explanation is granted for high-risk systems.

04 Privacy and data protection

Data minimisation, purpose limitation and least-privilege principles apply; anonymisation is mandatory if personal data is used in training data.

06 Security and robustness

Systems are tested against adversarial attacks, prompt injection, data poisoning and model theft; red-teaming is conducted for high-risk systems.

08 Social value and child protection

Systems targeting individuals under 18 apply age-appropriate language, avoidance of emotional manipulation, a ban on profiling and parental consent.

GOOD PRACTICE EXAMPLE

In CV pre-screening, an AI score cannot eliminate any candidate without review by an authorised person. This is the human oversight principle in practice.

06 · AI GOVERNANCE STRUCTURE

AI Governance Structure

AI projects, applications and infrastructure needs are managed through clearly defined roles. At the centre of this structure is the AI Governance Committee, positioned under the Digital Transformation Committee.

AI Governance Committee

The AI Governance Committee is one of the central elements of the framework that we have already implemented. It is the decision-making body that sets the strategic direction of AI use across the organisation, oversees risks and brings different functions together at one table. Thanks to its multidisciplinary structure, the perspectives of technology, legal, risk, data protection and business units converge at a single point of governance.

At a minimum, the following functions are represented on the Committee:

Technology & Digital Business Development	AI and Technology	Risk, Compliance & Internal Audit
Information Security	Legal	Data Protection
Human Resources	Strategy and Finance	Business Sectors

Key roles and responsibilities of the Committee

Strategy and direction

Sets the organisation's AI strategy and roadmap, defines priorities and updates them as needed.

Risk and approval decisions

Defines a governance approach that reflects the corporate risk appetite; makes the final decision on deploying high-risk systems and critical projects.

Approved tool management

Approves the list of AI tools authorised for use across the organisation; evaluates requests for new tools and licences.

Principles and ethics oversight

Oversees the application of responsible and ethical AI principles; reviews and redirects uses that conflict with them.

Incident management and response

Provides direction during critical AI incidents; exercises kill-switch authority in emergencies.

Monitoring and reporting

Tracks governance KPIs, submits regular status reports to senior management and drives continuous improvement of the framework.

Guidance and exception management

Resolves ambiguities arising in the interpretation of policies and evaluates exceptions.

06 · AI GOVERNANCE STRUCTURE

Decision-making mechanism

The Committee meets regularly and holds extraordinary sessions for critical cases or incidents. Decisions are taken by majority; for high-risk systems, the concurrence of the risk and audit function is sought, and unresolved matters are escalated to senior management bodies. All decisions are minuted with their rationale.

Roles and Responsibilities Matrix (RACI)

Responsibilities for core processes are assigned using the RACI method. This clarity aims to prevent gaps and overlaps in accountability.

GUIDING PRINCIPLE

Every process has one, and only one, accountable party.

STAKEHOLDER	CORE RESPONSIBILITY
AI Governance Committee	Sets the strategy and roadmap, prioritises critical projects, approves the tools list.
IT Service Operations	Manages the technical infrastructure, implements data loss prevention systems, prepares the proposed list of secure tools.
Internal Audit, Control and Risk	Participates in risk assessments, conducts audits, investigates alleged ethical breaches.
Legal	Reviews contracts for IP and competition law implications; provides legal opinions.
Data Protection	Monitors data protection regulations, manages privacy and contractual processes.
Human Resources	Organises mandatory AI training within the corporate learning programme.
Sector IT Representatives	Submits project requests through the defined process and ensures integration.

WHEN ARE A FRIA AND A DPIA MANDATORY?

A Fundamental Rights Impact Assessment (FRIA) is mandatory for all high-risk systems; for systems that determine opportunities, access or rights for individuals; and for systems targeting under-18s that involve biometric data, emotion analysis or profiling. A Data Protection Impact Assessment (DPIA) is required for systems that process special categories of data, operate at large scale, perform systematic monitoring or transfer personal data abroad. Development cannot begin until these assessments are complete.

07 · AI PROJECT LIFECYCLE

AI Project Lifecycle

AI, advanced analytics and machine learning projects are run using a standard, traceable and measurable method. Projects draw on established data science methodologies such as CRISP-DM, TDSP and OSEMN. The lifecycle consists of seven stages, each ending with a defined decision point.

1

Ideation and preliminary assessment

The problem statement, target output, expected benefit and project class are defined; the risk assessment form is completed and submitted for approval.

2

Project initiation and conceptual design

A project manager is appointed; the business problem is translated into an analytical problem; scope, success criteria and business objectives are defined.

3

Data collection and preparation

Data sources, access rights and data quality are analysed; licence terms of external data and open-source components are reviewed.

4

Development and modelling

The solution architecture is defined, different models are compared, performance is evaluated and versions are recorded.

5

Testing and model evaluation

Model performance is tested, results are validated with the business unit, and go-live criteria are assessed.

6

Deployment and monitoring

The solution goes live and integrations are completed; data drift and model performance are analysed, and business objectives are tracked.

7

Project closure

Final documentation is completed, lessons learned are recorded, the benefit measurement plan is activated and results are communicated transparently across the organisation.

Intellectual property and financial governance

Rights to code, designs, text and inventions produced by employees using approved tools are set out clearly and in writing under the applicable legislation, notably FSEK and SMK.

On the financial side, costs such as cloud, model training, licences, APIs and token consumption must be transparent and traceable. A cost owner is assigned to every project, and lower-cost alternatives delivering the same business value are regularly evaluated.

08 · DATA SECURITY AND CONTROL APPROACH

Data Security and Control Approach

Corporate data is divided into four classes: Public, Internal, Confidential and Restricted. For each class, the appropriate use in each AI environment is clearly defined. This classification provides an actionable framework for which data may be processed in which environment.

DATA CLASS	APPROVED ENTERPRISE AI	GENERAL / PUBLIC AI	IN-HOUSE PRIVATE LLM / RAG
Public	Permitted	Permitted	Permitted
Internal	Permitted	Controlled	Permitted
Confidential	Approved tenant only	Prohibited	Controlled
Restricted	Special approval and isolation	Prohibited	Isolated environment only

Input and output security

Requests to AI systems pass through layered controls: personal and sensitive data scanning, content filtering and prompt injection prevention. On the output side, hallucination detection, prevention of personal data and trade secret leakage, and content labelling are applied. Chat interfaces disclose up front that content is AI-generated.

Data loss prevention (DLP)

Data loss prevention for AI tools is applied in layers: controls on uploads from user devices, network traffic and content semantics work together. For highly sensitive data types such as ID numbers and IBANs, blocking and incident logging are triggered automatically. Semantic DLP goes beyond classic keyword matching to evaluate the context of a sentence.

Particular care with employee data

Emotion recognition and biometric categorisation in the workplace are prohibited under Article 5 of the AI Act. Employee monitoring systems are not deployed without consulting HR, Legal and employee representatives. Decisions on promotion, performance, disciplinary action or termination cannot be based solely on AI output.

09 · AGENT AND MCP GOVERNANCE

Agent and MCP Governance

The spread of agent-based applications has made the governance of connectivity infrastructure a policy area in its own right.

The framework addresses this area early on as a dedicated policy topic.

MCP Gateway architecture

All MCP traffic is routed through a central gateway; direct connections are not permitted. The gateway is expected to provide authentication via the corporate identity infrastructure, fine-grained authorisation, DLP integration, rate limiting and comprehensive audit logging. The authority to halt traffic instantly in an emergency rests with the Governance Committee.

Agent control matrix

Agents operate within clearly defined boundaries:

CONTROL	RULE
Scope	The scope, permissions and intended use of each agent are clearly defined and approved.
Autonomous actions	Actions that modify data require human approval for a defined observation period. After this period, only low-risk actions are opened up to automation.
Data access	Agents may access only registered, classified and authorised data sources, and only via the central gateway.
Usage quota	Periodic token limits apply: an alert at a defined threshold, an automatic stop at the upper limit.
Sandbox	New or updated agent configurations are tested in a sandbox before going into production.

GOOD PRACTICE EXAMPLE

Defining an observation period for autonomous agents and requiring human approval for every critical step during that period preserves control while capturing the benefits of automation. Token quotas and automatic stops provide protection in terms of both cost and security.

10 • EMPLOYEE AWARENESS AND SHADOW AI PREVENTION

Employee Awareness and Shadow AI Prevention

However strong the technical controls, a single misinformed employee transferring data to a personal tool can undermine the entire structure.

That is why training and awareness are an integral and ongoing part of the framework.

Three-tier training approach

01 All employees

AI awareness, acceptable use, shadow AI risk and personal data protection fundamentals.

02 Approved tool users

Tool-specific use, prompt security, data classes and incident reporting. Before first use, then annually.

03 Developers and data scientists

OWASP LLM Top 10, MCP security, modelling and RAG; maintained annually and with updates.

KEY PRINCIPLE

The impact of training should be measured by behavioural change, not completion rates.

AI Champions Network

Selected employees from each department and subsidiary are appointed as AI Champions. Champions are responsible for local tool support, case reviews and awareness activities, and share knowledge through quarterly meetings. The network acts as a bridge between central governance and practice on the ground.

Preventing shadow AI

The use of unapproved tools on employees' own initiative is defined as shadow AI and poses a serious risk at enterprise scale. Our prevention strategy has two axes: on the supply side, the approved tool catalogue is expanded and business needs are met quickly; on the demand side, training and clear guidance are provided. When such use is detected, the priority is awareness and meeting the legitimate need rather than sanctions; if the need is legitimate, the tool is referred to the Committee for review.

11 · INCIDENT RESPONSE AND BUSINESS CONTINUITY

Incident Response & Business Continuity

AI incident response plan

Sudden degradation in model performance, hallucination, prompt injection, risk of unauthorised data leakage, an increase in discriminatory output, discovery of shadow AI and an agent exceeding its authorised scope are all treated as AI incidents. Response follows eight steps:

1 Detection and logging	2 Classification	3 Isolation	4 Shutdown if critical
5 Root cause analysis	6 Corrective action	7 Verification	8 Reporting

No system is restored until testing and remediation are complete.

Notification timelines

Defined notification timelines apply to personal data breaches, critical incidents in high-risk systems, and high-impact prompt injection or model theft. Since a single incident may fall into more than one category, obligations to notify data subjects, regulators, customers under contract and internal stakeholders are assessed separately.

Business continuity

AI systems, particularly in agent-based use, are not put into production without outage planning. A dependency map is maintained for every critical system, and recovery time and acceptable data loss objectives are defined. To reduce vendor dependency, alternatives are evaluated for critical services and failover is tested through regular drills.

KEY PRINCIPLE

An architecture locked into a single provider is not a technical choice but a strategic risk. Data portability is secured by contract, and failover to alternatives is tested through regular drills.

12 · AUDIT, TARGETS AND CONTINUOUS IMPROVEMENT

Audit, Targets and Continuous Improvement

The effectiveness of governance is tracked through defined indicators, which make the framework's progress in practice measurable.

#	GOVERNANCE INDICATOR
01	Registration in the AI inventory
02	FRIA completion rate and model card availability for high-risk systems
03	Mandatory training completion rate
04	Share of suppliers with contractual AI clauses in place
05	Share of incidents handled within SLA and compliance with critical incident notification
06	Average closure time of internal audit findings and number of reported ethical breaches

Reporting structure

Quarterly

AI Officer → AI Governance Committee



Semi-annual

Governance Committee → Digital Transformation Committee



Annual

Digital Transformation Committee → Board of Directors

The AI Officer reports quarterly to the AI Governance Committee; the Committee submits a semi-annual status report to the Digital Transformation Committee, which presents an annual effectiveness and compliance review to the Board. This tiered reporting makes governance visible up to the highest level of management.

Document management and transition

The document is reviewed at least annually; regulatory changes, critical incidents or business model changes trigger an extraordinary review. A defined transition period and compliance screening plan bring AI systems already in use into line with the framework, making it a living guide rather than a static document.

GLOSSARY

Glossary

Brief explanations of the main terms and abbreviations used in this document are provided below.

TERM	DESCRIPTION
EU AI Act	EU Artificial Intelligence Act. The EU regulation governing AI systems according to their risk level.
KVKK / GDPR	Turkish (Law No. 6698) and European Union regulations on the protection of personal data.
SMK / FSEK	Industrial Property Law (6769) and Law on Intellectual and Artistic Works (5846); IP legislation.
OECD / UNESCO Principles	Internationally recognised principles and recommendation frameworks for responsible and ethical AI.
OWASP LLM Top 10	A security reference defining the most critical security risks in large language model applications.
LLM	Large Language Model. A type of AI model that understands and generates text.
MCP	Model Context Protocol. A protocol that enables AI models to connect securely to external data sources and tools.
MCP Gateway	The central gateway through which all MCP traffic passes for authentication, authorisation and auditing.
Agent	An AI system that can plan steps and take actions using tools in pursuit of an assigned goal.
RAG	Retrieval-Augmented Generation. A method that grounds the model's response in information retrieved from corporate data sources.
FRIA	Fundamental Rights Impact Assessment. An assessment of the impact of high-risk systems on fundamental rights.
DPIA	Data Protection Impact Assessment. Analyses the privacy risk of personal data processing.
RACI	A method for assigning Responsible, Accountable, Consulted and Informed roles.
DLP	Data Loss Prevention. A control layer that prevents sensitive data from leaving the organisation without authorisation.
Hallucination	When an AI system confidently produces information that is not actually true.
Prompt injection	A type of attack that uses malicious input to alter a model's instructions.
Kill-switch	An emergency control that instantly halts an AI system or its traffic in critical situations.
Sandbox	A secure, isolated environment where new configurations are tested before going into production.
Data drift	The deviation of data distribution or model performance from initial conditions over time; monitored regularly in live systems.
Model card	Standardised documentation describing a model's purpose, limitations, performance and risks.
CRISP-DM / TDSP / OSEMN	Established project and process methodologies used in data science and analytics projects.
Shadow AI	The use of unapproved AI tools on employees' own initiative, outside corporate approval.

CLOSING NOTE

This document provides a general framework for corporate AI governance. Organisations should adapt it to their own risk profiles and the regulations that apply to them.

At Zorlu Holding, we regard this framework as a living guide that evolves with our maturity and is continuously updated in line with technological, regulatory and ethical developments in AI; we are committed to keeping it current. We hope it will inspire organisations and professionals across the field.