


Practical Manual for AI Implementation in Enterprises

Guidelines and Best Practices



AI

Bratislava University of Economics and Business
Faculty of Business Management

Bratislava, 2026

Practical Manual for AI Implementation in Enterprises

Guidelines and Best Practices

Bratislava University of Economics and Business, Faculty of Business Management
Dolnozemska cesta 1, 852 35 Bratislava, Slovak Republic

Project Information

This manual was developed as part of the outputs of the AI-ImpactSK project, No. 09I05-03-V02-00003/2025/VA

The project is funded by the European Union - NextGenerationEU through the Recovery and Resilience Plan of the Slovak Republic.

Prepared within the AI-ImpactSK project

Copyright

© 2026 All rights reserved.

No part of this publication may be reproduced, stored or transmitted in any form or by any means without the prior written consent of the AI-ImpactSK project team.

Bratislava, 2026

Contents

About This Manual	4
How the Process Works.....	5
How to Use the Manual by Stage of AI Adoption.....	6
STEP 1 WHAT BUSINESS PROBLEM ARE WE SOLVING?	7
STEP 2 WHERE ARE WE NOW?	8
STEP 3 WHAT SHOULD WE DO?	10
STEP 4 IS THE SOLUTION FEASIBLE?	11
STEP 5 HOW WILL WE TEST IT?	12
STEP 6 HOW WILL WE KNOW WHETHER IT WORKS?	13
STEP 7 WHAT COULD GO WRONG?	14
STEP 8 WHAT MAY EMPLOYEES DO?	15
STEP 9 HOW WILL WE RUN THE PILOT?	16
STEP 10 WHAT HAPPENS AFTER THE PILOT?	17
FINAL RECAP	20

About This Manual

This manual was developed within the AI-ImpactSK project to translate research findings into a practical process that enterprises can use when deciding how to introduce and further develop artificial intelligence. It is intended for enterprise management, process owners and teams responsible for information technology, data, security, compliance and human resources.

The manual draws on the empirical and analytical outputs of the AI-ImpactSK project. Its evidence base includes three separate quantitative surveys conducted within two waves of data collection, in-depth interviews with enterprises, expert webinars, segmentation, statistical and econometric analyses, comparative assessments, and relevant regulatory, ethical and methodological frameworks. Their synthesis made it possible to identify recurring conditions, barriers, risks and best practices associated with AI implementation in business practice.

The key findings are translated into a ten-step decision process. It guides an enterprise from defining a specific problem through assessing readiness, selecting a suitable use case and designing a pilot, to measuring results, managing risks, setting rules for use and deciding on the next course of action. Each step ends with a decision gate that helps determine whether the enterprise can proceed, whether the solution needs to be modified, or whether implementation should be paused or terminated. The process is therefore sequential, but not strictly linear.

How to Use This Manual

The manual is practical and methodological in nature. It does not replace legal, security, sector-specific or certification assessment of a particular AI system. The scope of controls and documentation should be adapted to the specific use case, the type of data, the degree of autonomy of the solution and the associated risks.

How the Process Works

Before moving through the individual steps, it is useful to understand the logic of the overall process. This section explains how decision gates, required outputs and responsibilities are linked so that implementation can be managed continuously and adjusted as needed.

Unified Implementation Map

The map connects ten steps into a single decision process. Each step begins with a specific management question, produces a required output and ends with a gate at which the enterprise decides whether it can proceed. The process also allows the solution to be modified or paused, or for the enterprise to return to an earlier step if new gaps, risks or information emerge during implementation.

Table 1: Unified Map of the Ten-Step AI Implementation Process

Step	Guiding question	Required output	Gate
1	What business problem are we solving?	Problem definition + designated owner	G1
2	Where are we now?	Readiness profile + gap-remediation plan	G2
3	What should we do?	Prioritised use case	G3
4	Is the solution feasible?	Data, technology, legal requirements and risk-based governance regime	G4
5	How will we test it?	Pilot card + value mechanism	G5
6	How will we know whether it works?	Baseline + multi-level KPIs	G6
7	What could go wrong?	Risk register + controls	G7
8	What may employees do?	Minimum AI use policy + training	G8
9	How will we run the pilot?	Indicative plan + evidence	G9
10	What happens next?	SCALE / MODIFY / EXTEND / TERMINATE	G10

CORE PRINCIPLE: Problem first, technology second. Every tool, data source and control mechanism is assessed in relation to a specific process and expected business value.

LIFE CYCLE: ADOPTION = start of use; **IMPLEMENTATION** = introduction into a task or process; **INTEGRATION** = incorporation into systems, responsibilities and measurement; **SCALING** = controlled expansion of a validated solution.

Roles and Decision Rights

At the beginning of the process, executive sponsorship, process ownership, technical implementation, control and approval must be clearly separated so that every decision and risk has a specific owner (responsible person). In a smaller enterprise, one person may hold several roles, but explicit accountability and independent assessment of critical issues must not be lost. Combining management support, clear process ownership and user involvement helps prevent AI from remaining an isolated initiative without clear accountability.

Table 2: Roles and Decision Rights in the AI Implementation Process

Role	Primary responsibility
Executive sponsor	Approves resources, accepts residual risk and decides after pilot evaluation.
Process owner	Defines the problem, baseline, target state and operating conditions.
AI solution owner	Coordinates design, testing, documentation and the vendor.
IT / data / security	Verify architecture, data quality, access, event logging and cyber risks.
Legal / compliance / data protection	Assess the legal basis, personal data, transparency and relevant obligations.
Users and output reviewer	Use the solution, verify outputs and report incidents.

DECISION RIGHT: Each gate must have a designated approver. The person performing a control must have the authority to pause the project if a legal, security or ethical condition is not met.

How to Use the Manual by Stage of AI Adoption

The ten-step manual can be used at every stage of AI adoption, but the way it is applied and the scope of individual assessments will differ. An enterprise with no prior experience can use it as a complete process for preparing its first pilot, while an advanced enterprise can use it mainly as a recurring decision and control mechanism for new or materially changed use cases. The stage of adoption therefore does not determine which control requirements may be omitted, but where the enterprise should start and which steps require greater attention.

Table 3: Applying the Manual by Stage of AI Adoption

Stage of AI adoption	How the enterprise should use the manual	Key steps	Outcome
Enterprise without an AI adoption plan	First use the manual to identify a suitable business problem and verify basic readiness. Proceed to pilot selection only after Gate G2 has been passed.	Step 1: problem definition; Step 2: readiness assessment; once readiness conditions are met, continue with Steps 3-10.	Decision on whether the enterprise is ready to start a controlled low-risk pilot or must first address identified gaps.
Enterprise planning AI adoption	Apply the full Steps 1-10 process to the first specific use case. Place particular emphasis on selecting the right pilot and verifying its feasibility.	Especially Steps 2-5: readiness, prioritisation, feasibility and the pilot card; then Steps 6-10 during the pilot.	A prepared and approved pilot with a defined objective, responsibilities, data, KPIs, risks and decision rules.

Stage of AI adoption	How the enterprise should use the manual	Key steps	Outcome
Enterprise in the pilot stage	If a pilot is already under way, retrospectively verify that Steps 1-5 were properly completed. Add any missing evidence and continue managing the pilot through Steps 6-10.	Especially Step 6: KPIs; Step 7: risks; Step 9: indicative 90-day pilot plan; Step 10: final decision.	Evidence-based decision to SCALE / MODIFY / EXTEND / TERMINATE.
Enterprise in an early stage of AI use	Use the manual as a common gate-based process for each new use case and also to review existing solutions. Readiness is no longer assessed only at enterprise level, but also in relation to the specific process.	Step 2 periodically; Step 3 at portfolio level; especially Steps 4, 6, 7, 8 and 10 for each solution.	Transition from isolated experiments to a governed portfolio of validated AI solutions with owners, KPIs and controls.
Advanced AI user	The manual is not a one-off guide to starting adoption, but a recurring portfolio governance system. Every new or materially changed model, vendor, data source or form of use must pass through the relevant gates again.	Step 2 after a significant change or periodic reassessment; Step 3 for portfolio prioritisation; Steps 4, 6-8 and 10 for integration, performance, risk and AI governance; Steps 5 and 9 for a new or higher-risk pilot.	Regular decisions on scaling, modifying, replacing or terminating solutions and control of costs, models, data and vendor dependencies.

Regardless of the current stage of adoption, the basic logic remains the same: every new or materially changed use case must pass through the decision gates that are relevant to its scope and risk.

STEP 1 WHAT BUSINESS PROBLEM ARE WE SOLVING?

Defining the Business Problem

The first step turns a general interest in AI into a specific and measurable business problem, identifies the affected process and its owner, and compares AI with a simpler alternative. Implementation therefore starts with a problem or opportunity, not with selecting a tool. Expectations of greater efficiency or time savings should therefore be linked from the outset to a baseline, clear process ownership and a method for measuring the result.

Step Output

The output of this step is a concise business record containing:

- a description of the process, problem or untapped opportunity and identification of the process owner;
- the affected groups and a baseline for time, cost, quality or risk;
- the expected measurable change and target state;

- a justification for the suitability of AI and comparison with a simpler non-AI alternative;
- the date by which a decision on the next course of action will be made.

Control Questions

- Is the problem or opportunity significant enough to justify pilot testing?
- Can the expected value be measured against a baseline?
- Is AI proportionate, or could the problem be solved by changing the process, setting a rule or using conventional automation?
- Is there a designated process owner who will accept the outcome and be accountable for the change?

GATE G1

CONTINUE: The problem or opportunity is specific, measurable and business-relevant, has a designated owner, and there is a justification for testing AI.

DO NOT CONTINUE: An unclear problem, missing owner or inability to establish a baseline means returning to process definition.

STEP 2 WHERE ARE WE NOW?

Diagnostic Assessment of Enterprise Readiness

In the second step, the enterprise assesses whether it has the basic strategic, organisational, data, technological, human and governance conditions required for controlled AI implementation. Each area is evaluated using a diagnostic question and a score from 1 to 5. The assessment should be carried out jointly by enterprise management, the process owner and, as needed, representatives of IT, data, security or legal functions. The result is an indicative readiness profile and a list of gaps for which appropriate measures must be defined before the pilot begins.

Table 4: Diagnostic Assessment of Enterprise Readiness for AI Implementation

Area	Diagnostic question	Score (1-5)	Enterprise priority
Strategy	Does the enterprise have a clearly defined objective for AI implementation and is this objective linked to business priorities?		
Management support	Does senior or middle management support the planned implementation and is a specific project sponsor designated?		
Data	Does the enterprise have available, high-quality, current, representative and legally usable data required for implementation?		

Area	Diagnostic question	Score (1-5)	Enterprise priority
Technology and integration	Is the existing technology infrastructure compatible with the planned AI solution and does it allow secure integration?		
Capabilities	Does the enterprise have internal or accessible external expertise for selecting, testing, controlling and operating the AI solution?		
Processes	Is the selected process sufficiently documented, repeatable and measurable, and does it have a designated owner?		
Ethics and human oversight	Does the enterprise have rules for responsible AI use, output verification, human oversight, escalation and safe shutdown?		
Privacy and security	Have data protection requirements, access rights, cybersecurity and event logging been assessed?		
Funding	Are investment and operating costs known, including licences, data, integration, training, human review and monitoring?		
Outcome measurement	Are baselines, target KPIs, data sources and responsibility for measurement defined before implementation or the pilot?		

Assessment method: 1 - not met: the condition has not been established or addressed; 2 - initial: only informal or partial measures exist; 3 - partially implemented: the condition is met for the pilot or a limited area; 4 - implemented: the condition is documented and applied consistently; 5 - integrated: the condition is systematically managed, monitored and improved.

The score reflects the enterprise's current state, not the desired or planned state.

The average score is calculated as the total points divided by ten:

1.00-2.49 - low readiness: the enterprise should not yet begin a complex AI project,

2.50-3.49 - transitional readiness: a clearly defined low-risk pilot is possible,

3.50-5.00 - high readiness: the enterprise has good conditions for preparing and conducting a controlled pilot.

CRITICAL CONDITION: The overall average score is only indicative and cannot compensate for a fundamental gap in safe and responsible implementation. The pilot should not proceed unless at least Level 3 is achieved for legality and usability of the required data, privacy and security, clear accountability, appropriate human oversight, and the ability to stop the solution safely. These conditions apply regardless of the priority assigned by the enterprise.

GATE G2

The enterprise has reached at least transitional readiness, meets the critical conditions, and has a specific measure, responsible person and resolution deadline for all other identified gaps. Readiness of the specific use case is then verified in Steps 3 and 4.

STEP 3 WHAT SHOULD WE DO?

Use-Case Prioritisation Matrix

The third step turns technology-driven experimentation into a transparent selection of a business-relevant use case. Candidate use cases are compared in terms of expected value, feasibility, risk, measurability and the suitability of a non-AI alternative, so that the first selected project offers real value within a controllable scope. The use case should therefore be assessed in relation to the specific process, available data, capabilities, expected value and the ability to measure the outcome reliably.

Table 5: AI Use-Case Prioritisation Matrix

Criterion	Question for the specific use case	Score
Business relevance	Does it address a significant problem or support an important business priority?	1-5
Economic value	Is the net benefit significant and measurable?	1-5
Non-AI alternative	Is the use of AI more appropriate than a simpler solution?	1-5
Frequency and reach	How often will AI be used and by how many users?	1-5
Data readiness	Are the data available, high quality and usable?	1-5
Technical feasibility	Is integration realistic within the planned time horizon?	1-5
User readiness	Will users accept the solution and be able to review its outputs?	1-5
Risk	How severe would an incorrect output be?	1-5; reverse-scored
Pilot measurability	Can a baseline and comparison logic be established?	1-5

Note: Prepare this table separately for each use case.

SCORING: For all criteria except **Risk**, a higher score indicates a more favourable assessment. For **Risk**, reverse scoring is used: **1 = high or difficult-to-control risk**, **5 = low or well-controlled risk**. This ensures that, across all criteria, a higher score indicates greater suitability of the use case for a pilot. This reverse scoring is used only for use-case prioritisation. In the risk register in **Step 7**, likelihood and impact are scored in the standard direction, where a higher score indicates higher risk.

STEP OUTPUT: A prioritised list of candidate use cases, with a brief justification of their value, feasibility, measurability and controllability of risks, and selection of one use case for the first pilot.

PREFERRED FIRST PILOT: High expected business value + good technical and data feasibility + controllable risk + an available baseline and measurable outcome. Do not choose a critical process with potentially irreversible impact as the first pilot.

GATE G3

One priority use case has been selected and the reasons for its selection and for not selecting the alternatives are documented.

STEP 4 IS THE SOLUTION FEASIBLE?

Data, Technology, Legal Requirements and a Risk-Based Governance Regime

The fourth step verifies whether the selected use case can be tested safely, technically and lawfully. It assesses data and technology readiness, data protection and security, model characteristics, vendor conditions, regulatory requirements and operational arrangements. Based on the results, the enterprise determines an appropriate intensity of governance and control.

Simplified regime is suitable for low-impact cases that do not involve personal, sensitive or regulated data and do not include autonomous interventions in business processes.

Standard regime applies to ordinary internal implementation with a controllable impact and appropriate human oversight.

Enhanced regime is intended for cases with a higher impact, sensitive or regulated data, significant legal or security requirements, or a higher degree of solution autonomy. In a smaller enterprise, one person may hold several roles, but the basic control functions must not disappear.

The simplified, standard and enhanced regimes are internal governance categories used in this manual. They do not replace the legal classification of an AI system or expert assessment of obligations arising from applicable regulation.

Table 6: Minimum Feasibility Assessment of the AI Solution

Area	Minimum verification
Data	Source, owner, quality, currency, representativeness, legal usability.
Personal data protection	Purpose, minimisation, legal basis, retention period, impact assessment where needed.
Technology	Integration, availability, latency, limits, monitoring, rollback capability.
Security	Access controls, encryption, environment separation, event logs, misuse testing.
Model	Purpose, version, known limitations, accuracy, hallucinations, bias, explainability.
Vendor	Contracts, subcontractors, processing location, use of inputs for training, exit plan.
Regulation	Assessment of the regulatory regime of the use case, transparency, human oversight, documentation, and need for further legal or expert assessment.
Operations	Support, incidents, model change, continuity, costs as usage grows.

STOP CRITERIA: Unclear legal basis, unauthorised or unacceptable processing of personal, sensitive or regulated data, inability to provide appropriate human intervention where required, uncontrollable impact, insufficient security, or contractually unacceptable use of data.

Feasibility Decision

Open conditions should be converted into specific measures, responsible persons should be assigned, and the necessary approvals should be obtained. The pilot should not proceed if a critical data, legal, security or organisational gap remains without an appropriate measure, responsible person or resolution deadline.

STEP OUTPUT: The enterprise prepares a concise record containing open conditions and responsible persons, controls required before the pilot, accepted limitations, the designated risk-based governance regime, approving persons and deadlines.

GATE G4

A risk-based governance regime has been determined and the pilot can be carried out within the defined scope without unacceptable legal, security, ethical or operational risk.

STEP 5 HOW WILL WE TEST IT?

One-Page Pilot Card

The one-page pilot card brings together the hypothesis, value creation mechanism, scope, KPIs, human oversight, costs and stop criteria. Its purpose is to link technical validation of the solution to a specific process or organisational change and to define in advance what evidence will be required for the next decision. The card captures the basic pilot design; measurement of results, risk management and rules of use are developed in greater detail in Steps 6-8.

STEP OUTPUT: The enterprise prepares a one-page pilot card containing: pilot name, sponsor and risk-based governance regime; process owner and AI solution owner; a hypothesis in the form If..., then..., because...; value mechanism AI -> task -> process -> outcome; required process or organisational change; scope - included / excluded; users and method for comparing results; inputs, outputs and system links; preliminary KPIs and method for obtaining the baseline; human oversight - who reviews what, when and according to which criteria; duration, milestones and success conditions; estimated total costs and immediate stop criteria.

EXAMPLE OF A GOOD HYPOTHESIS: If employees use an assistant to produce the first draft of a standard response, median processing time will decrease by at least 25%, while factual error rates will not exceed the baseline and every response will be approved by a person before being sent.

GATE G5

The pilot card has been approved by the sponsor and process owner and, according to the designated governance regime, by the relevant technical, security, legal or control functions. The pilot scope, responsibilities, preliminary success criteria, human oversight and stop criteria are clearly defined.

STEP 6 HOW WILL WE KNOW WHETHER IT WORKS?

KPI Overview and Value Measurement

The sixth step creates a predefined evidence base for deciding on the next course of action and separates perceived benefit from objectively measurable process, quality, economic, user, risk and organisational outcomes. Greater speed or frequency of use alone, without control of quality and actual value, does not constitute evidence of success. Project analyses showed that the financial effects of AI may be heterogeneous and that short-term changes cannot automatically be interpreted as a consequence of AI implementation. The next decision should therefore be based on a combination of several types of evidence.

Table 7: KPI Overview and Value Measurement for AI Implementation

Type of evidence	Examples of KPIs	Recommended source and measurement frequency
Perceived	Satisfaction, perceived time savings	Short user survey - at defined milestones.
Process	Time, throughput, productivity	Process records - continuously; comparison with baseline.
Quality	Error rate, corrections, complaints	Sample review of outputs - weekly.
Economic	Net savings, revenue, avoided loss	Cost and performance data - monthly or by stage.
Use and adoption	Share of suitable tasks, frequency, extent of corrections	Usage and workflow records - weekly.
Risk	Incidents and residual severity	Incident register - continuously; formal review at the gate.
Organisational	Standardisation, capabilities, acceptance	Process audit and feedback - after each stage.

Minimum Evaluation Logic

- Use the same indicator definition before and during the pilot and, where possible, use a control group or a comparable period.
- Total costs include licences, data, integration, security, legal assessment, training, human review, corrections, monitoring and solution termination.
- Indicative net value = demonstrated benefits - total cost of ownership - costs associated with risks and remediation.
- A scaling decision must not rely solely on subjective assessment when objective data can be obtained.
- A short-term pilot does not automatically demonstrate long-term effectiveness or improved profitability.

GATE G6

Before the pilot starts, KPI definitions, baselines and target values, data sources, measurement frequency, responsible persons and success thresholds have been approved.

STEP 7 WHAT COULD GO WRONG?

Risk Register and Control Measures

The seventh step ensures that, before the pilot, the enterprise identifies significant technical, data, legal, human and organisational risks, assigns owners and defines appropriate preventive, detective and corrective controls. Assessment continues during the pilot because risks change and may accumulate as integration deepens. The risk register is therefore not a one-off document, but a continuously updated pilot management tool.

Table 8: Examples of Typical Risks and Control Measures

Typical risk	Recommended control or response	Responsibility
Incorrect or fabricated output	Human review, source checking and an escalation threshold	Process owner / reviewer
Leakage of confidential data	Data classification, leakage prevention and access control	Data owner / security
Discrimination or bias	Testing of relevant groups and escalation of deviations	Risk owner
Cyber misuse	Adversarial test scenarios and tool restrictions	Security
Over-reliance on automation	Fallback procedure and prohibition of autonomous critical actions	Process owner
Vendor dependency or model change	Exit plan, version control and monitoring of changes	AI solution owner
Low adoption or circumvention	Training, feedback and use audit	Process owner
Unplanned costs	Budget limits and continuous cost monitoring	Sponsor / finance function
Poor-quality or outdated data	Quality checks, source versioning, update rules and verification of data used	Data owner
Unclear accountability	Designation of the process owner, output approver and escalation rules	Management / process owner

For each specific pilot, the enterprise records at minimum the risk, its likelihood and impact, the owner, the control measure, residual risk after the control is implemented, the deadline or resolution status, and the escalation route.

Assessment

Likelihood and impact are both assessed on a 1-5 scale, where 1 is the lowest and 5 the highest level. Risk score = likelihood × impact. The enterprise defines in advance thresholds for acceptance, mitigation, escalation and unacceptable risk. After controls are implemented, the risk is reassessed as residual risk; its acceptance must be approved by the designated person.

INCIDENT: In the event of a serious incident, stop the affected function, preserve records of the event, limit its impact, inform designated responsible persons, perform a root-cause analysis and restart the pilot only after the corrective action has been approved.

GATE G7

Each significant risk has an assigned owner, likelihood and impact assessment, appropriate control measure, assessed residual risk, resolution deadline or status, and an escalation mechanism. Unacceptable residual risk prevents the pilot from continuing.

STEP 8 WHAT MAY EMPLOYEES DO?

Minimum AI Use Policy and User Readiness

The eighth step translates general principles into clear boundaries for everyday use, human review and escalation. Employees must know which tools and data they may use, what requires approval, what is prohibited and how to respond to an error, ambiguity or incident. The purpose of the minimum policy is to establish clear boundaries within which AI can be used and further developed without individual experimentation remaining outside enterprise governance.

The following rules provide a basic indicative framework that the enterprise should adapt to the tools, data, processes and risk profile of specific use cases.

Table 9: Minimum Rules for AI Use

Permitted	Only with approval	Prohibited
Work with public or approved internal data in an authorised tool	Personal, confidential or regulated data only in an approved tool and regime and subject to the applicable conditions	Enter passwords, access credentials or other security secrets into an AI tool; enter non-public data into an unapproved tool
Drafting, summarisation and analytical support with human review	Automated decisions, external outputs or system actions	Present an unverified output as a fact or final professional decision
Prompt reporting of an error, incident or ambiguity in accordance with the established procedure	A new model, integration, extension, agentic system or vendor	Bypass event logging, security controls or human oversight

Minimum Policy Content

- Scope, approved tools and use cases, prohibited data, and rules for approving new tools, use cases or significant changes.
- Verification of accuracy, sources, hallucinations, bias, copyright and output suitability.
- Specific oversight regime: a person approves every output, a person monitors continuously, a person performs system-level oversight, or autonomous action is prohibited.
- For agentic systems: permitted tools, permissions, transaction limits, confirmation of irreversible actions, emergency stop and complete logs.
- Protection against malicious prompt injection, rules for extensions, external integrations and publication of outputs.
- Incident reporting, transparency, accountability and regular policy review, including assignment of responsibility for updates.

AI LITERACY: Training must match the specific role. A regular user needs to know the rules for safe use and output verification; the process owner needs to understand value measurement and outcome control; the technical team needs to understand model limitations, monitoring and security; control functions need to understand audit-trail, compliance and escalation requirements.

GATE G8

Users involved in the pilot have been trained appropriately for their role, have confirmed that they are familiar with the rules for AI use, have access to a support channel, and know the procedure for stopping the solution and reporting an incident.

STEP 9 HOW WILL WE RUN THE PILOT?

Indicative 90-Day Pilot Implementation Plan

The ninth step links pilot preparation and execution in an indicative 90-day schedule that progresses from alignment and controlled design through launch, verification and stabilisation to a formal decision. The objective is not full AI integration within 90 days, but to obtain sufficient and comparable evidence without endless experimentation. Project findings support gradual capability building while recognising the different complexity of enterprises and use cases.

Table 10: Indicative 90-Day AI Pilot Implementation Plan

Period	Objective	Key activities	Evidence / output
Days 0-15	Alignment	Problem definition, readiness, use-case selection, sponsor and team	G1-G3 approved
Days 16-30	Design and control	Feasibility, pilot card, KPIs, risk register, policy and training	G4-G8 approved
Days 31-45	Controlled launch	Small group, test data, monitoring, daily feedback	Stable pilot without a critical incident

Period	Objective	Key activities	Evidence / output
Days 46-60	Verification	Expansion within the approved scope, quality, costs, system resilience test scenarios	Interim assessment of KPIs and risks
Days 61-75	Stabilisation	Adjustments to processes, training, controls and technical settings	Change log and solution version
Days 76-90	Evaluation	Analysis of benefits, costs, risks, adoption and scalability	Decision package

Note: The 90-day horizon may be adjusted to the specific conditions and complexity of implementation.

Governance Rhythm

The pilot needs a governance rhythm appropriate to its stage so that incidents, changes in quality, costs and user barriers are captured in time and ongoing adjustments remain traceable. The difference between declared use and systematic implementation is also associated with the absence of regular monitoring, ownership and decision points.

Table 11: Governance Rhythm for Pilot Implementation

Frequency	Content	Participants
Daily at launch	Incidents, obstacles, output quality	Pilot team
Weekly	KPIs, risks, costs, feedback and changes	Process owner + control functions
30/60/90 days	Formal decision gate	Sponsor + steering group
EVIDENCE DISCIPLINE: Retain model and prompt versions, dates of changes, approvals, output samples, control results, incidents, costs and feedback. Without this audit trail, the results are not reliable.		
HORIZON FLEXIBILITY: The period may be shortened for a simple low-risk use case or extended due to data preparation, integration, regulation, seasonality or the need for a larger sample. It is not a deadline for full AI integration.		
GATE G9 The pilot was conducted within the approved scope and the decision package contains sufficient, comparable and auditable evidence.		

STEP 10 WHAT HAPPENS AFTER THE PILOT?

SCALE / MODIFY / EXTEND / TERMINATE

The tenth step converts pilot results into a transparent decision on the next course of action and limits the influence of technical attractiveness or sunk costs. The enterprise compares results with predefined thresholds for value, quality, risk, adoption and operational readiness. The heterogeneous

effects identified in the research do not support automatic scaling; some solutions require modification, longer observation or termination.

Table 12: Decision Matrix After Pilot Completion

Decision	When it is appropriate	Required next step
SCALE	Value and quality targets have been met; risk is controlled; adoption and operations are stable.	Scaling plan, capacity, governance, budget, monitoring and renewed approval of the scope.
MODIFY	There is a remediable gap in the process, data, solution, adoption or control.	Define the change, establish success criteria in advance for retesting, and conduct time-bounded retesting.
EXTEND	Results are promising, but the sample, duration or seasonality do not provide sufficient evidence.	Extend observation without a fundamental change; identify the missing evidence and set an end date.
TERMINATE	The benefit is insufficient, costs are disproportionate, risk is unacceptable, or another solution is better.	Terminate safely, manage data and access, and document lessons learned.

Decision Card

Document the evidence, conditions and individual accountability for the decision. The research showed the need to link perceived benefits with objective indicators, full costs and conditions of implementation maturity.

Step Output

- the selected decision and the date on which it was made;
- evidence on value, quality, adoption and risk;
- total cost of ownership and comparison with expected benefit;
- assessment of readiness for operations and support;
- conditions, measures, deadlines, the approver of the final decision and a description of accepted residual risk.

NOT ACCEPTABLE: Do not scale simply because the pilot was technically interesting, users rated it positively, or costs have already been incurred. The decision must be based on predefined success criteria, KPIs and risk thresholds.

AFTER THE DECISION TO SCALE

Transition from Pilot to Operations

After a positive decision to scale, the enterprise should verify that pilot results will remain sustainable at greater scale and in normal operations. Scaling is not simply an increase in the number of licences because it changes data volumes, the number of users, costs, the consequences of errors, and requirements for support and governance. The research interprets AI as a progressively developed organisational capability whose value depends on complementary data, process, human and governance assets.

Table 13: Checklist for the Transition from Pilot to Operations

Area	Control question when scaling
Value	Will the unit benefit be maintained at greater volume and across different user groups?
Process	Is the new workflow standardised, with clear accountability and a fallback procedure?
Technology	Are capacity, availability, integrations, monitoring and recovery adequate for operations?
Data	Have data quality, provenance traceability, access, retention and changes in data distribution been addressed?
People	Are training, support, change communication and a capability plan in place?
Governance	Is a model or solution owner designated, with change approval and periodic reassessment?
Risk	Have risks been reassessed for the new scope and affected groups?
Vendor	Are service levels, costs, exit arrangements, model changes and subcontractors contractually managed?

MINIMUM OPERATING STANDARD: Performance and drift monitoring, assessment of alternatives, audit records, incident management, regular access reviews, testing after changes and the ability to roll back safely.

REASSESSMENT: Triggered by a significant change in the model, task, data, purpose, users, integration, vendor, legal conditions, costs or risk. After 3, 6 and 12 months, or at another interval proportionate to the nature and risk of the solution, decide CONTINUE / OPTIMISE / REASSESS / RETIRE.

GATE G10

The selected decision - SCALE / MODIFY / EXTEND / TERMINATE - has been formally approved and has defined conditions, responsible persons and a deadline for the next step. If the decision is to SCALE, the process owner has accepted the solution into operations and the sponsor has approved the transition to operations based on an updated business case and risk assessment.

FINAL RECAP

The final overview allows management, before the final decision, to verify the completeness of the process and traceability of evidence across all decision gates. It is not a new form, but a concise recap of what should already have been produced in the individual steps. The project synthesis shows that success does not depend on a single dimension, but on the combination of value, readiness, capabilities, governance, risk and measurement.

Table 14: Summary Checklist of the Ten Decision Gates

Gate	Control statement
G1	The business problem or opportunity is specific, measurable and business-relevant, has a designated owner, and there is a justification for using AI.
G2	Critical readiness conditions are met and measures, responsible persons and resolution deadlines are defined for all other identified gaps.
G3	A priority use case has been selected and the reasons for its selection and for not selecting the alternatives are documented.
G4	A risk-based governance regime has been determined and the pilot can be carried out within the defined scope without unacceptable legal, security, ethical or operational risk.
G5	The pilot card is approved and clearly defines scope, responsibilities, preliminary success criteria, human oversight and stop criteria.
G6	KPI definitions, baselines and target values, data sources, measurement frequency, responsible persons and success thresholds are defined.
G7	Significant risks have assigned owners, assessments, control measures, assessed residual risk and an escalation mechanism.
G8	Users have been trained appropriately for their role, know the rules for AI use, have access to support and know the procedure for incidents or stopping the solution.
G9	The pilot was conducted within the approved scope and the decision package contains sufficient, comparable and auditable evidence.
G10	The decision to SCALE / MODIFY / EXTEND / TERMINATE has been formally approved and has defined conditions, responsible persons and a deadline for the next step.

If all gates have been demonstrably closed, the enterprise formally confirms one of four next-step options - scale, modify, extend or terminate - and assigns conditions, deadlines, responsible persons and the date of the next review to the decision. Transparent ownership and periodic reassessment support the transition from individual experimentation to governed AI integration.

AI-ImpactSK

Practical insights for responsible, value-oriented implementation of artificial intelligence in enterprises

ai-impact.sk

Bratislava University of Economics and Business, Faculty of Business Management

Project AI-ImpactSK, No. 09I05-03-V02-00003/2025/VA, is funded by the European Union - NextGenerationEU through the Recovery and Resilience Plan of the Slovak Republic.

Bratislava, 2026