

AI / Digitalisation

AI is becoming easier to access. Converting it into measurable enterprise performance is not.



AI-generated illustration
for explanatory purposes.

INSIDE THIS EXECUTIVE INSIGHT

EXECUTIVE PREMISE

AI creates lasting value when management redesigns economically important workflows, builds the complementary capabilities and governs verified results rather than tool adoption.

CONTENTS

Executive Lead Why now: the constraint has moved	03
1. AI amplifies the operating model it enters	04
2. The unit of value is the workflow, not the tool	05
3. Choose where AI earns the right to scale	06
4. Redesign before you automate	08
5. Govern economics, not AI activity	09
6. A focused 90-day start	10
7. The strategic decision for owners and CEOs	12
Sources	13
Dr. David Sven	

EXECUTIVE LEAD | WHY NOW

Executive Lead | Why now: the constraint has moved

For owners and chief executives, the artificial intelligence discussion has already moved through curiosity, experimentation and a proliferation of licences, copilots, chatbots and isolated use cases. The next stage is harder. Technology access is no longer the only binding constraint. The constraint increasingly sits inside the enterprise: process clarity, usable data, decision rights, management ownership, implementation capacity and control.

The executive decision is which economically important workflows deserve bounded redesign and pilot capital. Management should establish the baseline, prove that usable data and implementation capacity exist, assign a business owner and agree how finance will validate results. Scale follows evidence of operating improvement, full-lifecycle economics and effective controls. The policy opportunity is real, but technology access and adoption alone do not establish enterprise performance.

Why now: policy and capability context

This matters particularly in Vietnam now. Decision 1671/QĐ-TTg, approved on 28 August 2026, shifts national policy from research, development and application towards broader transformation through AI. It promotes redesign of management, production, business and services, treating people, data and infrastructure as foundations. Its 2030 ambitions are policy objectives, not achieved enterprise outcomes. Their management relevance is directional: the external operating environment is moving from experimentation towards scaled application. [1, 1a]

The legal environment has moved at the same time. Vietnam's Law on Artificial Intelligence, Law No. 134/2025/QH15, has been in force since 1 March 2026. Decree No. 142/2026/ND-CP, effective 1 May 2026, details important implementation mechanics, and Decision No. 33/2026/QĐ-TTg, effective 15 August 2026, establishes the current list of high-risk AI systems. The Law classifies systems by risk and differentiates responsibilities by legal role; it also provides transition periods for systems already in operation. The management implication is not that every corporate use case carries the same compliance burden. It is that legal role, risk classification, human oversight, documentation and other applicable controls must enter the design early. [2, 2a, 3, 4]

The strategic risk is therefore not simply late adoption. It is failure to build the complementary capability required to turn AI into operating performance. Being early is not itself an advantage: model quality and economics are moving quickly, so later adopters may gain access to stronger technology at lower cost and avoid some stranded experiments. But AI adoption and productivity are closely tied to human capital, digital capabilities and enabling technologies. A later technology purchase does not remove the need to build those complements. The advantage of bounded early action arises only when experimentation is converted into reusable process knowledge, data discipline, governance routines and implementation experience. [5]

MANAGEMENT IMPLICATION

The management question is no longer: Where can we use AI? It is: Which few workflows can materially change the economics of our company if we redesign them now, and who is accountable for proving it?

OPERATING MODEL

1. AI amplifies the operating model it enters

AI does not enter a vacuum. It enters an operating system that already contains the company's real decision logic: who does what, which information is trusted, where approvals sit, how exceptions are handled, what gets measured and how quickly problems are corrected. Where AI becomes materially embedded in a workflow, those characteristics shape the result. Acceleration is not automatically improvement.

In a disciplined workflow with clear ownership, stable data definitions and explicit controls, AI can reduce cycle time, improve consistency, expand capacity and support faster decisions. In a fragmented workflow with duplicated handoffs, weak master data, unclear authority and informal exception management, the same technology can also transmit inconsistency: errors move faster, outputs become harder to reconcile, shadow processes multiply and management gains activity without necessarily gaining control. The effect is contingent, not universal. Task design, technology quality and human capability matter, and narrow assistive applications can create direct task-level productivity gains even without enterprise-wide redesign. The stronger claim is therefore one of complementarity, not technological determinism.

That is why the sequence 'buy technology, train users, search for benefits' is structurally weak when the objective is material enterprise performance. It treats the tool as the transformation and the organisation as a passive recipient. Narrow assistive use cases can create useful value with limited process change. But when AI is expected to alter workflow economics, decision quality or control, the stronger sequence starts with the business constraint: identify where performance is limited, determine whether AI can change that constraint, and redesign the operating logic around the intended result.

The principle is therefore more precise than 'AI amplifies everything'. AI does not replace the operating model. Where it becomes materially embedded, it tends to increase the value of strong process and management disciplines while exposing the economic cost of weak ones.

ECONOMIC UNIT

2. The unit of value is the workflow, not the tool

A serious AI programme needs a unit of economic analysis. The tool is usually the wrong unit. Licence counts, user activity and prompt volumes tell management whether technology is being used; they do not tell an owner whether the company is becoming more productive, responsive or valuable. The economically relevant unit is the workflow in which time, labour, decisions, inventory, service and risk are converted into an output.

That workflow may be sales quotation, procurement approval, demand planning, production scheduling, maintenance triage, customer-service resolution, credit review, month-end closing or management reporting. The label matters less than the baseline. Before AI is introduced, management should understand current elapsed time, labour intensity, rework, error frequency, decision latency, conversion, service level, working-capital effect, control burden and management attention.

Only then can the business formulate a value hypothesis. If the objective is faster quotation turnaround, the baseline and target must be explicit. If the objective is to release engineering capacity, the constrained hours must be visible. If the objective is better inventory decisions, the mechanism must connect forecast quality to stock, availability, obsolescence and cash. Without that bridge, a technically successful demonstration can still be a commercially unsuccessful investment.

This distinction exposes a central digitalisation trap: adoption is not by itself evidence of transformation. High user activity can coexist with unchanged economics. Companies can complete successful proofs-of-concept but fail to scale them when the business case, integration path, data maturity or required organisational changes are not sufficiently resolved. A company can therefore deploy many AI tools and still retain the same approval chains, reconciliation effort, meeting load and decision delays. Technology has then been added, but the operating model has not materially changed. [6]

CAPITAL ALLOCATION

3. Choose where AI earns the right to scale

The next management problem is portfolio selection. Potential AI use cases can easily exceed the organisation's capacity to implement them properly. Capital matters, but so do senior-management attention, process-owner time, data-remediation capacity and the ability of control functions to support change without becoming a bottleneck themselves.

Four tests separate an investable workflow from an attractive demonstration. First, economic materiality: if this workflow improves, will the effect be visible in earnings, cash, capacity, service, risk or strategic speed after full-lifecycle cost? Second, absorption readiness: is the process sufficiently understood, is usable data available or realistically fixable, and can key dependencies be controlled? Third, decision and control feasibility: are accountability, human review, legal, cyber and model-risk requirements clear enough to support deployment? Fourth, management capacity and time to proof: can the organisation redesign, implement and measure the workflow without overloading the people whose attention is already critical elsewhere?

These tests deliberately combine opportunity with executability. A workflow with large theoretical value but no viable data path may not deserve first priority. A technically easy use case with negligible economic consequence may be useful for learning, but it should not be confused with strategic transformation. Full-lifecycle economics also matter: vendor and model costs, integration, data remediation, controls, recurring human review, change effort and potential switching costs all belong in the investment case. The objective is a small portfolio in which material economics and controllable execution overlap.

This is an owner-level decision because it allocates scarce organisational capacity. It also imposes discipline on vendors and internal enthusiasts: the most impressive demonstration does not automatically deserve capital. AI investment should begin where enterprise economics are material and absorption is controllable.

The four tests can be condensed into a capital-allocation screen. Economic materiality and absorption readiness determine the primary position; decision/control feasibility and management capacity remain mandatory funding gates.

AI WORKFLOW INVESTMENT SCREEN

Capital follows economic materiality and absorption readiness.

Technical novelty is not the test. Economic materiality runs from high (top) to low (bottom); absorption readiness runs from low (left) to high (right).

BUILD THE CONDITIONS		PRIORITY REDESIGN & PILOT	
High value, low readiness		High value, high readiness	
Fix process, data or control constraints before rollout.		Allocate bounded capital. Prove economics before scale.	
DEFER		LEARNING USE CASE	
Low value, low readiness		Low value, high readiness	
No strategic capital case. Avoid distraction.		Useful for capability building, not strategic transformation.	

Funding gates: decision and control feasibility; management capacity and time to proof.

Figure 1. AI Workflow Investment Screen | VIET TA synthesis

AI-generated diagram for explanatory purposes.

The upper-right quadrant is the natural candidate for bounded redesign and pilot capital. High materiality without readiness does not justify premature rollout; it identifies where management must first build the conditions for execution.

EXECUTION SEQUENCE

4. Redesign before you automate

Once a workflow has been selected for material economic impact, the task is not to insert AI into every existing step. That often automates the legacy architecture, including unnecessary approvals, duplicated checks and unclear responsibilities. This is not a rule that every low-risk assistive use requires process redesign. It is a rule for workflows in which AI is expected to change economics, decisions or control. The stronger question is what that workflow should look like if management designed it today around the required outcome, the available data and the appropriate level of human judgement.

Redesign therefore reaches beyond technology. Process steps may disappear or move. Data may need to be standardised at source rather than reconciled later. Human roles may shift from producing first drafts to validating exceptions, resolving ambiguity or making higher-consequence decisions. Decision rights must be explicit because an AI-supported recommendation creates no value when nobody is authorised to act on it. Controls belong inside the workflow, not as a layer added after launch. Technology architecture then supports the target design rather than defining it.

This is also where the workforce question becomes concrete. The relevant issue is not simply how many tasks AI can perform. It is how roles change when routine production, search, summarisation or pattern recognition becomes cheaper while judgement, accountability, exception handling and cross-functional coordination become relatively more important. Training should therefore follow the redesigned work. Generic AI literacy helps, but it does not substitute for role-specific capability inside a changed process.

Skipping redesign often creates a second operating layer. Employees use AI informally while the official process remains untouched. Outputs are copied back into old templates, checked manually, reconciled in spreadsheets and approved through unchanged hierarchies. The company then carries both the old process and the new technology cost. That is digital activity, not operating-model improvement.

GOVERNANCE

5. Govern economics, not AI activity

Governance becomes effective when it follows the economics of the workflow. The business owns the benefit. Technology and data functions enable the solution but cannot own the commercial or operational result on the business's behalf. Finance fixes the baseline, agrees the attribution method and validates realised gains. Where exact causality cannot be isolated, comparison periods, control groups where feasible or conservative attribution should be agreed before launch. Legal, risk and cyber functions define non-negotiable constraints, review requirements and escalation triggers. An implementation lead manages dependencies, adoption, training and exceptions across those functions.

The owner or board sponsor has a different role: set ambition, protect management capacity, define risk appetite and release capital by stage. That avoids two opposite failure modes: a central programme too broad to manage, and a collection of local pilots that no one stops because each looks inexpensive in isolation.

The funding logic should therefore resemble disciplined capital allocation. Establish a baseline. Define the target economic outcome. Set a bounded pilot budget and explicit stop/go criteria. Measure performance against the baseline, including full-lifecycle cost, control effectiveness and management burden. Do not count theoretical time savings as earnings unless capacity is actually redeployed, removed or converted into additional output. Then scale, redesign, defer or stop. A pilot that creates learning but fails the economic test can still be useful; it simply has not earned scale capital.

This discipline becomes more important under a dedicated AI regulatory framework. Under the current Vietnamese framework, the precise obligation set depends on legal role, risk classification, whether a system falls within the current high-risk list, transition status and the surrounding data, cybersecurity and sector regime. The Law and implementing instruments can make transparency, incident management, conformity assessment, technical records, human oversight and intervention relevant to particular systems. These requirements can change workflow design, cost, documentation and deployment timing. They therefore belong in the business case from the beginning rather than in a compliance review after launch. [2, 2a, 3, 4]

ACTION AGENDA

6. A focused 90-day start

A focused 90-day decision cycle is a useful management default for many companies, not a promise that every AI transformation can be completed in 90 days. Its purpose is to establish facts, make a bounded investment decision, control implementation and generate enough evidence for the next capital decision.

During the first 20 days, management identifies a limited set of candidate workflows and establishes the facts. The work is diagnostic, not promotional: current economics, cycle time, labour intensity, rework, data defects, decision bottlenecks, existing systems, legal and cyber constraints, and the management capacity available for implementation. The output is a candidate portfolio with usable baselines, not a catalogue of ideas.

During days 21 to 35, the company applies the four selection tests and chooses a small pilot portfolio. For each selected workflow, it defines target economics, target operating state, an accountable business owner, a funding limit, control requirements and explicit stop/go criteria. Management should be willing to reject technically interesting ideas that fail the economic or organisational threshold.

During days 36 to 70, the selected workflows are redesigned and implemented within bounded scope. Data is prepared, technology is configured, human review and exception handling are defined, and required legal, cyber and model controls are built in before volume scales. Training is specific to the changed work. For high-risk, heavily integrated or legacy-dependent use cases, this phase may end with a validated design, controlled test or approval-ready implementation rather than scaled live operation. The target is control and evidence, not cosmetic speed.

During days 71 to 90, the company compares results with the agreed baseline and attribution method. Finance tests the economics. The business owner tests operational performance. Risk functions test control effectiveness. Senior management then makes a capital decision: scale, redesign, defer or stop. Ninety days will not complete every transformation, especially where legacy architecture, major data remediation or regulatory approvals are involved. It should be enough to establish control over the initiative and decide whether the company has earned the right to scale.

The sequence below is therefore a management control cycle: each stage produces the evidence required for the next decision, and Day 90 ends with a capital decision rather than an automatic rollout.

90-DAY DECISION AND CONTROL CYCLE

Ninety days establish control and evidence, not necessarily transformation completion.

Days 1-20 Diagnose	Baseline economics Process and data facts Constraints and capacity	Candidate portfolio with usable baselines
Days 21-35 Decide	Apply four tests Select bounded pilots Owner, target and controls	Funded pilot set with stop/go criteria
Days 36-70 Redesign and pilot	Target workflow Data and controls Role-specific adoption	Controlled test, validated design or live pilot
Days 71-90 Prove and decide	Validate economics Validate controls Test management burden	Scale, redesign, defer or stop

Capital decision: scale only what has earned the right to scale.

For high-risk, heavily integrated or legacy-dependent use cases, Day 90 may mean an approval-ready or validated state, not scaled deployment.

Figure 2. 90-Day Decision & Control Cycle | VIET TA synthesis
AI-generated diagram for explanatory purposes.

STRATEGIC DECISION

7. The strategic decision for owners and CEOs

AI and digitalisation are moving from a specialist technology agenda into the operating architecture of competition. In Vietnam, the new national strategy and a dedicated, now operational AI legal framework are accelerating that shift. But the competitive outcome will not be determined by policy targets or technology availability alone. Productivity effects depend materially on complementary digital, human and organisational capability. The decisive work therefore remains inside companies, workflow by workflow.

Owners and CEOs therefore face a more precise decision than whether to 'do AI'. They need to decide which few workflows are economically important enough to redesign now, whether the organisation can absorb that redesign, who owns the benefit and what evidence will justify further capital. That decision is narrower than an enterprise-wide transformation slogan, but strategically more consequential because it builds the operating capability to scale later.

Selective early action can preserve strategic option value, but being early is not itself an advantage. Model quality and inference economics are changing rapidly, so later adopters may obtain stronger technology at lower cost, while some early experiments will become stranded. At the same time, advanced AI adoption still builds on enabling digital technologies, skills and organisational capabilities, limiting the idea that firms can simply skip the capability-building stage. The value of early action lies in organisational learning that survives technology cycles, not in commitment to an early vendor or model. Waiting preserves cash, but it also defers that learning; broad uncontrolled rollout spends capital without proving it. The stronger path is controlled concentration. [5, 6]

The management question is therefore no longer: Where can we use AI? It is: Which few workflows can materially change the economics of our company if we redesign them now, and who is accountable for proving it?

THE MANAGEMENT QUESTION

Which few workflows can materially change the economics of our company if we redesign them now, and who is accountable for proving it?

Sources

[1] Government of Vietnam. Decision 1671/QD-TTg, 28 August 2026; national AI strategy to 2030, vision to 2045. Issuance record and official strategy summary. [Original source](#)

[1a] Government News. National AI strategy: redesign of management, production, business and services; people, data and infrastructure as foundations. [Original source](#)

[2] National Assembly. Law 134/2025/QH15, 10 December 2025; effective 1 March 2026. Official issuance record. [Original source](#)

[2a] Government News. Official English reference translation of the AI Law. Articles 3, 9, 11-14 and 35 inform the discussion of roles, risk, controls and transition. The Vietnamese enacted text governs. [Original source](#)

[3] Government of Vietnam. Decree 142/2026/ND-CP, 30 April 2026; effective 1 May 2026. Implementation of the AI Law. [Original source](#)

[4] Prime Minister. Decision 33/2026/QD-TTg, 30 June 2026; effective 15 August 2026. List of high-risk AI systems. [Original source](#)

[5] OECD. Technology diffusion. Complementary skills and digital capabilities, and the selection of more productive firms into AI adoption. This supports complementarity, not a universal causal effect. [Original source](#)

[6] OECD/BCG/INSEAD (2025). The Adoption of Artificial Intelligence in Firms. Business-case, data and organisational barriers to scaling. Survey evidence concerns the G7 and Brazil, not a Vietnamese 2026 sample. [Original source](#)

Sources checked on 6 October 2026 (Bangkok). The investment screen and 90-day cycle are VIET TA management synthesis, not measured regulatory or research outcomes.

Prepared with AI assistance. Editorial responsibility: VIET TA.