

THE AI INSTITUTE / OPERATING INTELLIGENCE

EXECUTIVE SYNTHESIS · 05 OCT / 2026



Your AI Investments May Share One Demand Risk

Different projects can fail together when they rely on the same customers, adoption timetable or contract assumptions.

Boards, C-suite and senior management

Research period: March 2025 – October 2026

First published: 5 October 2026

Current to 5 October 2026

PUBLICATION RECORD

An Institute thesis and decision framework

Before approving another AI commitment, test whether apparently separate investments depend on the same unproven demand and adoption assumptions.

WHAT LEADERS SHOULD TAKE FROM THIS PAPER

- | | |
|----|---|
| 01 | Group commitments by the demand assumptions they share. |
| 02 | Separate technical progress from customers' ability and willingness to pay. |
| 03 | Test slower adoption against fixed costs and contract terms. |
| 04 | Assign one owner to the combined exposure, not just each project. |

Research method

Original Institute weekly executive synthesis. Official institutional summaries and a cross-country working paper establish the economic context; one explicitly labelled preprint illustrates a contract-specific technical constraint. The common-demand review and hypothetical example are Institute analysis, not a measured portfolio-risk model.

First published 5 October 2026 · Updated 5 October 2026 · Research period March 2025 – October 2026 · Version 1.0 · Suggested citation: The AI Institute, *Your AI Investments May Share One Demand Risk* (2026).

CONTENTS

Inside this paper

01	Find the assumption that joins the projects Board agenda
02	Supply is not the same as a market What changed
03	Trace the same customer through every business case The combined position
04	Efficiency does not cancel a commitment Economics and contracts
05	Test adoption where the customer operates Across markets
06	Give the combined exposure an owner This week’s decision
M	Method and limitations
R	References

How to use this brief

Start with the decision and use the framework to challenge the current plan. Figures retain their original populations and definitions. Institute analysis is labelled, and limitations are recorded before the references.

Find the assumption that joins the projects

A board can approve a computing contract, an AI-enabled product and a strategic partnership in separate meetings, with different sponsors and different business cases. Yet all three may depend on the same customers changing how they work within the same year. Separate approval papers do not make those bets independent. Before approving the next commitment, ask management to show where the underlying demand is shared.

The immediate prompt is a departmental paper listed by the IMF on 2 October. Its published summary links proposed AI-service hubs in the Middle East and nearby regions to adoption capacity in importing countries and cross-border coordination. It also identifies financial vulnerabilities if investment runs ahead of adoption. These are conditional regional assessments, not a prediction that a particular project will fail.¹

The Institute's conclusion is narrower than a verdict on the AI market: a company needs to understand its combined exposure to one adoption assumption. This week, the useful board decision is how much spending can become difficult to reverse before that assumption has been tested. The answer may be to accelerate, stage or redesign commitments. It should not be chosen simply because several projects have different names.

This is also a different question from whether each project has a plausible return. A portfolio can contain individually reasonable proposals whose disappointments arrive together. The missing information is the link between them: the same buyer budget, implementation partner, approval process, energy agreement or date by which customers are expected to start paying.

Supply is not the same as a market

The IMF paper is new to this week's evidence base. Earlier OECD work supplies a wider comparison: its March model of OECD and G20 economies examines how trade and differing adoption paths shape gains from AI. Cheaper imports can help slower adopters, but the model does not make domestic adoption irrelevant to competitiveness. It is scenario analysis, not an observed return for an enterprise project. ²

Together, these sources suggest a useful challenge to a business case: who must change behaviour for this investment to earn its return? A supplier's improved capacity can be real while a customer's purchasing decision is delayed. The delay might reflect integration work, unclear ownership, inadequate data, a limited budget or a service that has not yet proved useful. Each explanation requires different evidence.

Distinguish technical availability, operational adoption and paid demand. A model can be available without being authorised for a workflow. A workflow can be in regular use without producing incremental revenue or a cash saving. A paying customer can still reduce volume at renewal. Treating these stages as interchangeable makes an apparently precise forecast difficult to challenge.

Evidence of strong demand should change the decision too. Recurring use, independently measured outcomes, credible renewal behaviour and a customer's funded implementation plan can justify an earlier commitment. The purpose of the review is not to manufacture caution. It is to identify which assumptions have earned confidence and which remain an aspiration repeated across several spreadsheets.

Trace the same customer through every business case

Start with the material AI commitments already made or proposed. Include minimum cloud spend, equipment, long-term service agreements, product development, integration teams and partnership funding. Do not add unlike figures into a dramatic total. Separate cash already spent, unavoidable future payments and discretionary spending that management can still stop. These are different decisions, even when they support the same programme.

For each commitment, identify the customer group and the behaviour that generates value. Then connect proposals that depend on that same behaviour. If a new product needs banks to complete deployment reviews, the computing plan should not assume those banks are already generating production traffic. Nor should the partnership plan count those same deployments as an independent source of demand.

An illustrative company might reserve capacity, hire an implementation team and subsidise a distribution partner to serve one anticipated customer cohort. If onboarding takes longer, usage revenue arrives later while some costs continue. This is a hypothetical example, not a reported incident. Its value is that it exposes a timing dependency that three separate project reviews could miss.

Use a single demand record for the linked decisions. Name the assumption, its owner, the evidence supporting it, the next observable milestone and the commitments affected if it changes. Keep disagreement visible: sales may see a qualified opportunity while operations sees an unstaffed implementation. Reconciliation should clarify the difference rather than force both teams to report the same optimistic number.

EXHIBIT / FOUR QUESTIONS FOR ONE SHARED ASSUMPTION

01 Demand Which customer action must occur? Institute decision framework	02 Timing When does useful paid use begin? Institute decision framework	03 Exposure Which payments remain if use is delayed? Institute decision framework	04 Choice What can still be changed, and by whom? Institute decision framework
--	---	---	--

Efficiency does not cancel a commitment

Technical efficiency can reduce an input without reducing the bill in the same proportion. That distinction appears in current research as well as procurement practice. An October preprint, ePACT, examines AI serving under preset electricity commitments that charge for both overuse and underuse. Its experiments and simulations are bounded by those contract assumptions; they do not establish savings for every data centre.³

The practical question is not whether to adopt that research system. It is whether finance and engineering are optimising the same economic outcome. Before carrying an efficiency gain into the budget, determine which charge actually changes. A lower workload may help a usage-priced service while leaving a minimum payment unchanged. Different commercial terms can turn the same technical improvement into different financial results.

Review the agreement rather than relying on the product category. Identify volume floors, expiry dates, renewal provisions, cancellation rights, transfer restrictions and the treatment of unused capacity where they apply. Obtain competent legal and procurement review for the actual contract. Do not assume that a provision familiar in one market is enforceable, available or standard in another.

Then test a delay, not just a lower final market size. Ask what happens if customers eventually adopt but do so after a capacity reservation, staffing decision or partner payment becomes fixed. This separates a good long-term opportunity from a poor sequence of commitments. A staged agreement may preserve the opportunity, although flexibility has a price that should be compared with the exposure it reduces.

Test adoption where the customer operates

An export-oriented AI-service hub faces a different demand question from an enterprise buying AI for its own workforce. The former needs credible external customers and workable cross-border delivery; the latter needs a functioning internal workflow. The IMF's regional assessment should not be converted into a universal ranking of countries or a presumption that building local infrastructure automatically creates local demand. ¹

Europe's policy discussion adds another distinction. In September remarks, the IMF Managing Director separated the near-term building boom from longer-term productivity effects. Those effects operate on different timetables. A board should therefore avoid treating investment activity as evidence that the operational gains in its own European business have already arrived. The speech is a policy assessment, not a company-level performance study. ⁴

In Latin America, earlier IMF analysis highlights how informality and limited access to finance, skills and digital services can constrain adoption. That does not describe every firm or customer. It does mean that a forecast based on a large formal-sector client should not be applied without testing to smaller businesses with different resources. ⁵

For North American and Asia-Pacific operations, test the buyer population directly rather than borrowing the regional results above. For African and Middle Eastern markets, distinguish the institution hosting capacity from the organisations expected to use it. Across all markets, language, connectivity, procurement authority and implementation capacity are questions for local evidence. They are not reasons to assign an entire region a single adoption speed.

The missing evidence is often organisational rather than national: a budget holder, an approved use, staff able to implement it and a reason to renew. A local partner can help establish those facts, but should not be treated as proof of demand merely because it has signed a cooperation agreement. Keep partner incentives and customer evidence separate.

Give the combined exposure an owner

Ask the chief financial officer and the operating sponsor to bring one joined-up view to the next investment discussion. The technology leader should explain the capacity assumptions; commercial leaders should show the demand evidence; procurement should identify what can still change. The board needs a decision about the combined position, not another committee for every project.

Choose a few assumptions whose failure would affect several commitments. For each, compare the plan with a credible slower-adoption case and a stronger-demand case. Specify the evidence that would move the organisation between them. Avoid assigning reassuring probabilities without a defensible method. A clear description of what is unknown is more useful than an unsupported decimal.

Set the next commitment around an observable milestone where possible: paid production use by an identified customer group, a completed deployment review, demonstrated renewal or a verified operating outcome. A milestone should change the decision, not merely document activity. If the organisation would spend the same amount regardless of the result, it is not functioning as a decision gate.

Finally, document what would justify moving faster. A competitor's announcement alone may not, but customers committing funds and completing implementation could. Retain enough flexibility to respond to evidence in either direction. The objective is to avoid discovering that several apparently separate AI investments were all waiting for the same customer to become ready.

RESEARCH RECORD

Method and limitations

Method

Original Institute weekly executive synthesis. Official institutional summaries and a cross-country working paper establish the economic context; one explicitly labelled preprint illustrates a contract-specific technical constraint. The common-demand review and hypothetical example are Institute analysis, not a measured portfolio-risk model.

Limitations

The IMF regional paper was available through its indexed official summary; direct retrieval was blocked and the full paper was not inspected. No claim here relies on its unpublished tables or numerical estimates. OECD results are model scenarios, not realised enterprise returns. ePACT is a preprint with bounded energy and service assumptions, not verified field-wide savings. Regional analysis is not customer-level evidence. This is general management analysis, not investment or legal advice.



Sources and limitations

Source facts link to the original publication. Institute analysis reconciles definitions and turns the material into a management proposition. Frameworks are decision aids, not claims of universal causality. Read every important statistic with its population, date, definition and caveat.

Research cutoff: 5 October 2026. Review cycle: six months, or earlier after a material change in the underlying facts, standards or policy.

REFERENCES

Evidence behind the thesis

01 IMF: Harnessing Artificial Intelligence's Potential in the Middle East and Central Asia

Official indexed summary dated 2 October 2026, Departmental Paper 2026/020. Direct retrieval blocked; full paper not inspected. Authors' views; conditional regional assessment.

02 OECD: AI meets trade

Filippucci and colleagues, 18 March 2026, AI Papers 57. Multi-country, multi-sector model for OECD/G20 economies; scenarios, not observed business returns. No country-specific numerical result reproduced.

03 Peng and colleagues: ePACT

v1 submitted 1 October 2026; preprint, HTML methods and limitations inspected. H20/H200 serving experiments and trace simulations under positive two-sided energy-deviation costs; no general savings claim.

04 IMF: Europe and the Global AI Race

Kristalina Georgieva, page dated 21 September 2026. Policy remarks distinguishing near-term investment and longer-term productivity effects.

05 IMF: How AI Can Boost Productivity in Latin America

Bakker, Chen and Vasilyev, 20 March 2025. Older regional analysis of informality and adoption conditions, not a current adoption census.

THE AI INSTITUTE

Make the decision.
Then measure what changes.

Research · Readiness · Capability

theai.institute