

THE AI INSTITUTE / OPERATING INTELLIGENCE

EXECUTIVE DECISION BRIEF · 13 OCT / 2026



Size the AI Market Around Customers Who Can Use It

An adoption headline is not a sales forecast. Find the buyers who can implement, pay and renew.

Boards, C-suite and senior management

Research period: October 2026 – October 2026

First published: 13 October 2026

Current to 13 October 2026

PUBLICATION RECORD

An Institute thesis and decision framework

The useful AI market is not everyone reporting adoption. It is the customer group that can implement the offer, pay for the whole service and keep using it.



WHAT LEADERS SHOULD TAKE FROM THIS PAPER

- | | |
|----|--|
| 01 | Define the buyer and adoption threshold before using a market percentage. |
| 02 | Separate customers ready to buy from those needing substantial implementation support. |
| 03 | Price integration, training and review before treating revenue as attractive growth. |
| 04 | Test retention and paid expansion within each customer group, not just initial interest. |

Research method

Institute synthesis of a 12 October vendor-sponsored survey release, an explicitly older official establishment comparator, October regional economic analysis and two research preprints. Customer categories and commercial recommendations are Institute analysis. Research supports questions to test, not a numerical revenue forecast.

First published 13 October 2026 · Updated 13 October 2026 · Research period October 2026 – October 2026 · Version 1.0 · Suggested citation: The AI Institute, *Size the AI Market Around Customers Who Can Use It* (2026).

CONTENTS

Inside this paper

01	Count customers, not enthusiasm Markets & Value
02	Two adoption stories can both be true What changed
03	The implementation gap may be the product Where value accrues
04	Cheaper delivery still needs a customer test Research direction
05	Carry the method across borders, not the forecast Global application
06	Put one customer group into the forecast 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.

Count customers, not enthusiasm

A large AI adoption number can make a new offer look easy to sell. The missing step is the customer. Who has a problem worth paying to solve, the authority to buy, the information the service needs and the people to put it into daily use? A percentage cannot answer those questions on its own.

The immediate decision is where to put sales and delivery resources. A business can pursue buyers already equipped to implement its offer, sell a supported service to customers who need help, or fund that help itself. Those are different business models. Combining them into one market forecast hides both the opportunity and its cost.

The Institute's view is simple: build the forecast from customer groups you can describe and serve. Use broad adoption research to identify questions and promising areas, not to turn every organisation using AI into a likely buyer of your product.

Two adoption stories can both be true

On 12 October, Microsoft published Singapore findings from IDC research it commissioned. The page describes 200 local respondents and says 41% report AI-related revenue growth or new revenue streams. Its footnote identifies a July global survey of 5,204 respondents. These are reported outcomes, not audited profits or a measure of demand for a particular supplier.¹

An older official comparator asks a different question. Singapore's Ministry of Manpower reported on 30 April that 28.5% of private-sector establishments with at least ten employees had started adopting AI, including planning and piloting. Only 3.8% integrated it into core processes.²

These figures must not be subtracted or read as a change over time. The populations, definitions and periods differ. The vendor release does not provide enough detail to treat its respondents as a representative count of every business. A market forecast needs the population behind the percentage, not just the most encouraging number.

Ask the commercial team to define what adoption means in its own pipeline. An employee experimenting with a chatbot, a department paying for software and a company renewing an embedded service are different signals. Record them separately. A free trial may establish curiosity; a paid renewal after implementation establishes something much closer to durable demand.

The implementation gap may be the product

The World Bank's October Latin America and Caribbean update argues that the ability to absorb AI matters more than access to advanced models alone. It examines constraints involving workers, firms and government. This is regional economic analysis, not a supplier sales forecast.³

For a seller, the implication is practical. A customer may need cleaned records, a redesigned approval process, staff training and someone to resolve errors before a model becomes useful. That work can support a valuable service business. It can also consume the margin of a product sold as effortless software. The difference belongs in the price and delivery plan.

Consider an illustrative offer that helps firms prepare customer quotations. One buyer has standard product data and clear pricing authority. Another has inconsistent records and relies on an experienced employee to interpret every exception. They may share a sector and company size while requiring very different delivery effort. Do not give them the same implementation allowance by default.

Separate three groups: ready to implement, viable with paid support, and not yet economical to serve. These are proposed planning categories, not labels for a customer's worth. A business may choose to invest in the third group for strategic reasons, but that investment should be visible rather than disguised as low-cost acquisition.

Cheaper delivery still needs a customer test

New research offers reasons to keep testing delivery costs. A study revised on 9 October compares six multi-agent workflows with a single-agent reference across ten benchmarks using the same model. One workflow improves overall accuracy; five do worse. Some configurations improve a particular task while consuming more tokens and performing worse overall. It is a preprint, not a study of customer retention. ⁴

The commercial lesson is not to avoid agents. It is to price the system that solves the customer's actual job. A more elaborate architecture may be worthwhile for difficult cases and wasteful for routine ones. Keep the cost of review, support and failed attempts alongside model charges when deciding which customers and tasks to serve.

A separate new retail preprint explores how AI-generated product shortlists could change assortment decisions. It combines model rankings with synthetic customer preferences, so it does not measure real shopping revenue. It suggests a direction to watch: AI may change how customers discover offers as well as how suppliers deliver them. ⁵

Carry the method across borders, not the forecast

The World Bank's East Asia and Pacific update distinguishes the region's role in AI-related trade from the wider spread of AI use. Infrastructure activity and application demand are not interchangeable measures. A company selling into a data-centre supply chain faces a different market from one selling workflow software to ordinary firms.⁶

For North America and Europe, test the target sector's procurement, security and integration requirements rather than assuming a mature technology market makes implementation easy. Across Asia-Pacific, Latin America, Africa and the Middle East, check local-language performance, support capacity, connectivity and payment arrangements for the specific buyer. These are diligence questions, not a ranking of regions.

The evidence reviewed does not provide comparable conversion or renewal rates across those markets. Keep that gap explicit. Start with a narrowly defined customer group in each proposed market and measure its economics. Expand because the service works for that group, not because a result from another country has been renamed a global opportunity.

Put one customer group into the forecast

Ask the sales, finance and delivery leads to rebuild one AI revenue assumption together. Name the buyer, the job, the purchase decision and the conditions required to go live. Show paid conversion, implementation effort, continuing support and renewal evidence separately. Where evidence is missing, label the assumption and decide how to test it.

Fund the support that makes a promising group commercially viable. Stop counting unqualified interest as committed demand. Watch whether lower delivery costs open a new customer group, and whether discovery shifts toward AI intermediaries. The prize is not the largest headline market. It is a repeatable offer that customers can use and your business can afford to deliver.

EXHIBIT / THREE DIFFERENT COMMERCIAL PLANS

READY	SUPPORT	LATER
Implement now Buyer has the conditions to use the offer. Institute proposed planning categories	Price the help Include the work needed to obtain value. Institute proposed planning categories	Make investment explicit Do not hide enablement costs in growth. Institute proposed planning categories

RESEARCH RECORD

Method and limitations

Method

Institute synthesis of a 12 October vendor-sponsored survey release, an explicitly older official establishment comparator, October regional economic analysis and two research preprints. Customer categories and commercial recommendations are Institute analysis. Research supports questions to test, not a numerical revenue forecast.

Limitations

The survey populations are not directly comparable. Self-reported outcomes do not establish causation, audited profit or willingness to pay. Benchmark and synthetic-preference studies do not prove production returns. Comparable regional conversion and renewal data were not established.



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: 13 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	Microsoft / IDC: Singapore AI business-value findings 12 October 2026 release; July global survey cited. Vendor-sponsored, 200 local respondents; reported outcomes, not representative national prevalence or causal profit evidence.
02	Singapore Ministry of Manpower: AI adoption among firms 30 April 2026 release, older context. Private-sector establishments with at least ten employees. Planning and piloting count as starting adoption; not comparable with IDC maturity categories.
03	World Bank: Latin America and Caribbean economic update October 2026 report, disclosed 6 October. Regional absorptive-capacity analysis; not measured commercial conversion or supplier revenue.
04	Fu et al.: Do More Agents Help? Preprint first submitted 4 June; revised 9 October 2026. Six workflows versus a reference using GPT-4.1 across ten benchmarks; results depend on this design.
05	El Housni and Ibn Brahim: AI-Mediation in Retail Markets 9 October 2026 new preprint, visible in 12 October batch. Model-generated rankings combined with synthetic customer preferences; not observed retail revenue.
06	World Bank: East Asia and Pacific economic update October 2026 update. Regional trade, adoption and skills analysis; underlying series use different populations and dates, not one global demand measure.

THE AI INSTITUTE

Make the decision.
Then measure what changes.

Start with seven questions and a free readiness report. Then choose an optional 10–15 minute verbal assessment with Maya, our AI interviewer.

Start your free assessment: readiness.theai.institute

Research · Readiness · Capability

theai.institute