ETHORITY®

AI BUSINESS OPERATING SYSTEM

RESEARCH EDITION 2.0 · WAVE 1 · OCTOBER 2026

RESEARCH STUDY · GLOBAL EDITION 2026/27

Put AI to work.
The Good AI Prism
Global Market Study 2026/27

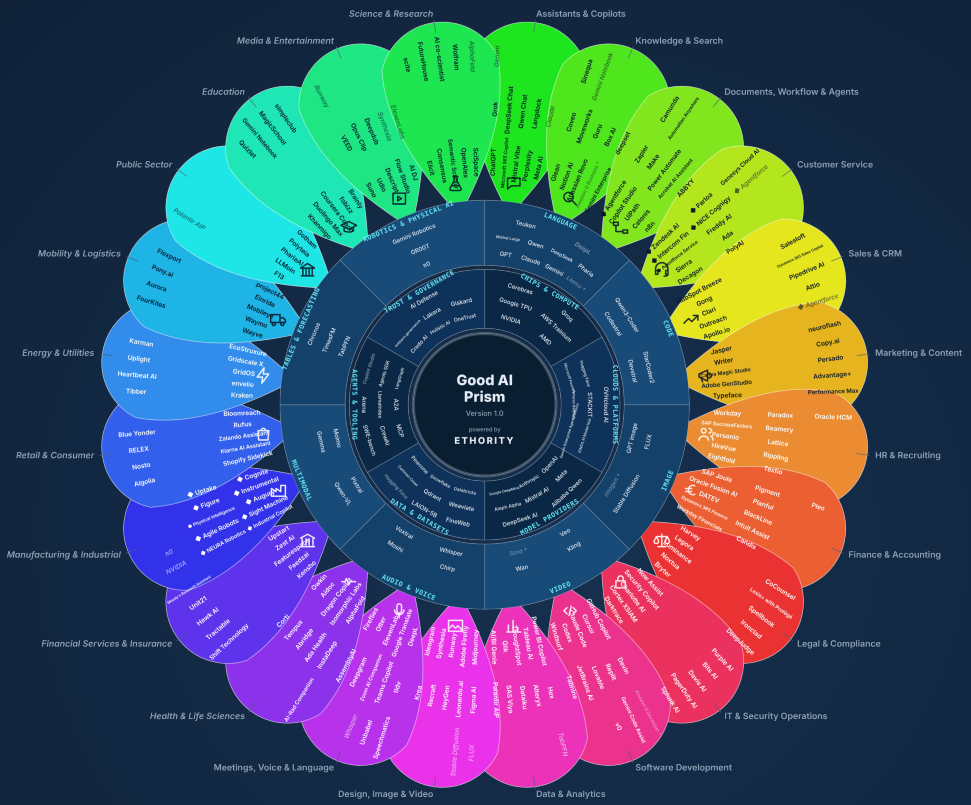
What businesses in Europe, the Americas, Asia-Pacific, the Middle East and Africa need from AI now, and the decisions that will shape 2027.

US$2.67 tn
AI SPENDING IN 2026

30+
ECONOMIES COMPARED

961
AI OFFERINGS MAPPED

10 of 107
DOCUMENT BOTH KEY DATA FACTS · WAVE 1



Created by **Sten Franke**
with a people and governance lens by Severin L. Wilson
October 2026

Independent market research
Global · regional · sector

IN 90 SECONDS

AI adoption is advancing unevenly. Organisations need practical ways to select useful applications, build capability and measure results. Worldwide AI spending is forecast to rise by half to US$2.67 trillion in 2026, and nearly nine in ten organisations in a global survey report regular AI use in at least one function. In the same survey, 44% report scaling it and 37% report any effect on profit; these are separate answers, not stages of one funnel. Official statistics make the gap plainer: in the EU, one firm in five with ten or more staff uses AI, and in the national studies we compared, large firms adopt more often than small ones, by 1.6 to 4.3 times. Know-how is the barrier named most often. Our review of 107 customer-service AI offerings shows how hard choosing is: we found public documentation of both an EU hosting option and limits on training with customer data for 10 of them: 8 state that customer data is not used for training and 2 use it only with opt-in consent. The finding is part of Wave 1 of ETHORITY's longitudinal research programme. Wave 1 documents the state of the provider information ETHORITY reviewed on 19–20 October 2026. Changes are recorded in subsequent waves.

STUDY ANSWER IN ONE TABLE

Five questions leaders are asking, and what the evidence says

QUESTION	FINDING	WHAT LEADERS SHOULD DO
How far has adoption come?	It differs widely. In a global survey about 90% of organisations report regular use and 44% report scaling it; statistics offices count between 5% and 42% of firms in organised use.	Move from tool trials to two or three measured use cases.
Where is the gap widest?	In small firms: EU 17% against 55% for large firms; Singapore 14.5% against 62.5%; South Korea 29% against 77%.	Offers and programmes built for small firms, including supplier programmes run by large ones.
What holds firms back?	Know-how is named most often (EU 71%, Germany 85%, 81% of Korean manufacturers without AI specialists). Legal uncertainty follows in Europe; in Brazilian industry, cost is named first.	Readiness and staff AI literacy first; a cost case per process.
Is regulation a blocker?	It is a set of dated deadlines that differ by jurisdiction: Korea and Vietnam now; US states 1 January 2027; EU high-risk duties 2 December 2027; India 13 May 2027.	One AI inventory, mapped to every jurisdiction you operate in.
Can buyers compare AI offerings?	Not easily. In 10 of 107 customer-service offerings we found both key data facts documented (8 no training, 2 opt-in only; Wave 1, provider information reviewed 19–20 Oct 2026); prices come in at least seven units.	Ask for documentation before demos; convert prices to your own volume.

The question is moving from whether to use AI to which AI, how safely, and whether it pays. The evidence shows how to answer it step by step: a useful task, a measured baseline, supported learning and a result that can be reviewed.

EXECUTIVE CONCLUSION

Ten findings

AI adoption is advancing unevenly. Organisations need practical ways to select useful applications, build capability and measure results. The ten findings below show where the gaps sit, with different emphasis in each region.

01 **Use is broad; scale is rare.**

Nearly nine in ten organisations in a global survey report regular AI use; 44% report scaling it and 37% attribute some EBIT effect to it (McKinsey, Aug 2026). Self-reported, separate answers, not a funnel.

02 **Official statistics tell a soberer story.**

Counted by statistics offices, one firm in five with ten or more staff used AI in the EU in 2025; about one in three in the UK; about one in eight Australian businesses.

03 **Large firms adopt more often than small ones in every study we compared.**

By 1.6 to 4.3 times in the selected national studies, whose size bands differ: EU 55% vs 17%, France 58 vs 15, Singapore 62.5 vs 14.5, UK 49 vs 28, Mexico 69.9 vs 43.7.

04 **Use by people and adoption by firms tell different stories.**

In France 47.8% of working-age people use generative AI (usage telemetry), while 18% of firms with ten or more staff use AI (official statistics). The measures differ; neither shows what happens inside a given firm.

05 **Know-how is the barrier named most often.**

Lack of expertise is named by 71% of EU firms that considered AI and 85% of German non-users; 81% of Korean manufacturers have no AI specialists. Legal uncertainty follows in Europe; cost leads in Brazilian industry.

06 **Spending grows by half, mostly into infrastructure.**

Worldwide AI spending is forecast at US$2.67 tn in 2026 (+49.5%), of which US$1.48 tn is infrastructure. Asia/Pacific spending grows 4.6-fold by 2030.

07 **AI has become a running cost.**

About one organisation in five says operating costs constrain use. In Germany licences are the smallest cost driver; cloud, data preparation and integration cost more. In Mexico, 46% struggle to understand AI cost structures.

08 **Agents are the growth story, and they widen the size gap.**

Spending on AI agents and assistants is forecast to grow from US$16.5 bn (2025) to US$65.5 bn (2027). 40% of large organisations scale agents, against 22% of smaller ones.

09 **Rules diverge, but the deadlines are concrete.**

AI laws are in force in South Korea and Vietnam; US state rules apply from 1 January 2027; EU high-risk duties from 2 December 2027; India's data-protection duties from 13 May 2027. Labelling of AI-generated content is the one rule most regimes share.

10 **Buyers can rarely compare offerings on facts.**

In 10 of 107 customer-service AI offerings we found public documentation of both an EU hosting option and limits on training with customer data (8 no training, 2 opt-in only; Wave 1, provider information reviewed 19–20 Oct 2026); prices come in at least seven different units.

© 2026 ETHORITY & STEN.AI LLC · Put AI to Work: The Good AI Prism Global Market Study 2026/27 · **ethority.ai/ai-prism-study**

CONTENTS

How this study is built, and where to start

Forty pages in five parts: the answer, the evidence, the market map, three worked cases, and the route from decision to operation. Every chapter separates what the evidence shows, what it means and what to decide now.

OVERVIEW

01 Cover and the Good AI Prism 1
02 Executive answer 2
03 Ten findings 3
04 Contents and reading paths 4
05 Founder's perspective 5
06 A second lens: work, people and oversight 6

THE MARKET MAP

15 Scope: the Good AI Prism 15
16 Taxonomy and units 16
17 What the market documents 17
18 Boundary conditions and coverage gaps 18

THE EVIDENCE

07 Global investment and spending 7
08 Adoption: two measures and the size gap 8
09 Europe and Germany 9
10 The Americas: United States 10
11 Asia-Pacific 11
12 Canada, Latin America, Middle East and Africa 12
13 Eleven sectors 13
14 Fourteen business functions 14

THREE WORKED CASES

19 Customer Service selection 19
21 Manufacturing quality support 21
23 Cross-border AI inventory 23

FROM DECISION TO OPERATION

25 Decision journey 25
26 Architecture and dependencies 26
27 Evidence record and controls 27
28 Full cost and accepted outcomes 28
29 Workforce and operating model 29
30 How ETHORITY helps 30
31 Capability status and proof 31
32 Roadmap: the first 90 days 32
33 Roadmap: 2027 33
34 Roadmap: evidence refresh 34
35 Scenarios and decision gates 35
36 Worksheet: decision and evidence 36
37 Worksheet: pilot and outcome 37
38 Method and limitations 38
39 Sources, credits and authors 39

Reading paths

READER	READ FIRST	THEN	USE
Boards and investors	Executive answer, ten findings	Full cost (28), scenarios (35)	Ask for accepted outcomes, not usage
Mittelstand and small firms	Adoption and size gap (8), Europe (9)	Case 1 (19–20), first 90 days (32)	Worksheet: decision and evidence (36)
Large and multinational enterprises	Boundary conditions (18)	Case 3 (23–24), architecture (26)	One AI inventory mapped to every jurisdiction
Public sector, science, education	Sectors (13), rules (18)	Evidence record (27)	Procurement-ready evidence, not demos
AI providers	What the market documents (17)	Evidence refresh (34)	Publish pricing unit, hosting and data-use terms
Media and analysts	Taxonomy and units (16)	Method (38)	Name the definition and size band of every figure

The evidence annex (linked at ethority.ai/ai-prism-study) carries the full country rankings, the regulation table by jurisdiction, Germany in depth, the ten offerings behind Exhibit 14, the risk surveys and the complete reference list of the first edition.

© 2026 ETHORITY & STEN.AI LLC · Put AI to Work: The Good AI Prism Global Market Study 2026/27 · **ethority.ai/ai-prism-study**

FOUNDER'S PERSPECTIVE

The decisions that will shape 2027

No business needs another reason to pay attention to AI. It needs a way to decide where AI belongs, what it must deliver, and who is responsible when it acts.

The figures in this study describe different parts of that challenge. Investment shows where capacity is being built. Adoption shows where organisations are experimenting. Neither tells a board which process to change, which supplier to choose, or whether the money will create value.

Our reading is that the advantage in 2027 will go to organisations that connect these questions. Start now: choose one important process, measure how it performs today, compare credible options, examine the evidence, and agree in advance when a pilot proceeds and when it stops.

> **"An AI decision should say what we expect to improve, what evidence supports the choice and who is accountable for the result. That is where putting AI to work begins."**
>
> Sten Franke, Founder, Principal Architect & Lead Creator

The Good AI Prism starts outside the organisation: what is available, what providers document and what is still unknown. A business decision adds what lies inside: objectives, data, systems, people, costs and responsibilities. ETHORITY connects the two, from market possibility to an accountable decision, and from a chosen pilot to an outcome you can measure.

Every one of those choices is now made under changing conditions: prices, providers, capacity, regulation and the organisation's own priorities. So we add one test to every AI decision: *will it still work if those conditions change?* Pages 26 and 35 show how to build that test into the choice.

This study is a starting point: a global view of adoption, a closer look at the evidence buyers need, and an agenda for the decisions to take next.



Sten Franke
Founder, Principal Architect & Lead Creator

Creator and architect of ETHORITY's AI systems, connecting business strategy, market intelligence and engineering.

HOW TO READ THIS STUDY

Each chapter separates three things: **what the evidence shows** (sourced), **what it means** (our interpretation) and **what to decide now** (a proposed action). The worksheets on pages 36 and 37 turn them into tools you can use on one decision.

A SECOND LENS · WORK, PEOPLE AND OVERSIGHT

AI changes work before it changes the organisation chart

A credible AI pilot needs a workforce plan alongside its technical plan. The evidence in this study shows why: in a Bitkom survey of German firms, 66% rate their staff's AI skills as low; in South Korea, 81% of surveyed manufacturers have no AI specialists; and where AI already touches decisions about people, the rules are explicit. Uses in recruitment and worker management are listed in Annex III of the EU AI Act, with duties from 2 December 2027, and Italy's national AI law requires employers to inform staff when AI is used at work.

The gains can be real. In a field experiment with about 70,000 applicants, AI-led interviews raised job starts by 18%. In customer support, measured productivity rose by 14% on average and by 34% for less experienced agents. But these gains depend on who reviews the outputs, how errors are handled and whether the released capacity is put to use.

The business case should therefore show more than hours saved. It should show whether the capacity released improves service, quality or growth, and who is responsible for that conversion.



Severin L. Wilson
Chief Business & Transformation Officer · Human Capital, AI Transformation & Governance

Connecting enterprise relationships, organisational transformation and responsible AI adoption.

Five questions to answer for every pilot

1. **Which tasks change?**
   List the tasks the AI takes over, supports or creates, before the pilot starts.
2. **Which skills do people need?**
   Name the skills for operating, reviewing and improving the system, and who provides the training.
3. **Who reviews consequential decisions?**
   Assign a named person for outputs that affect customers, applicants or employees.
4. **How can people challenge an error?**
   Provide a route for employees and affected people to question an output and get a human answer.
5. **What happens to released capacity?**
   Decide in advance whether freed time goes into service, quality or growth, and who is responsible for making that happen.

Sources: Bitkom, Sep 2026; KCCI via JoongAng Daily; Regulation (EU) 2026/1744; Norton Rose Fulbright on Italian law 132/2025; Erasmus field experiment; NBER working paper 31161. Page 29 turns these questions into roles and responsibilities.

1 · THE GLOBAL MARKET

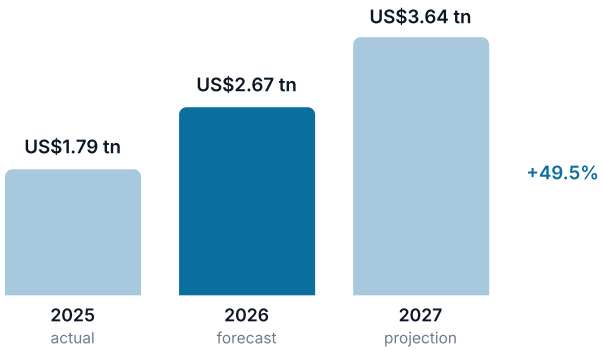
Big spend, broad use, rare scale

AI spending is forecast to rise by half in 2026, and more than half of it goes into infrastructure. In the wider economy, use is broad and shallow: most organisations run AI somewhere, and few run it at scale or can show a return.

EXHIBIT 1

AI spending grows by half in one year

Worldwide AI spending, US$ trillion



Source: Gartner, Sep 2026. Broad definition including infrastructure; 2027 is a projection.

© 2026 ETHORITY & STEN.AI LLC · Good AI Prism Study 2026/27 · ethority.ai/ai-prism-study

≈90%
of organisations in a global survey report regular AI use in at least one function

44%
report scaling it across the enterprise

37%
attribute some EBIT effect to AI

18.8%
of working-age people worldwide use generative AI

Sources: McKinsey, Aug 2026 (1,719 respondents, 97 countries; self-reported); Microsoft, Sep 2026 (telemetry, Q2 2026). The three survey figures are separate answers, not stages of a funnel.

EXHIBIT 2

Every region grows by a third or more a year

Latest regional AI spending forecasts

REGION	LATEST FORECAST	WHAT DRIVES IT	SOURCE
World (broad definition)	US$2.67 tn in 2026, +49.5%; US$3.64 tn in 2027	Includes infrastructure; agents and assistants US$29.2 bn (≈+77%, calculated)	Gartner, Sep 2026
United States	More than 55% of the global AI market in 2025	US and China together about 70%	IDC, Sep 2025
Europe	US$290 bn by 2029; 33.7% a year	Software 58.5% and banking 12.5% of 2026 spend	IDC, Apr 2026
Asia/Pacific	US$121.5 bn (2025) to US$555.2 bn (2030); 35.5% a year	China about half; Japan fastest at 39.8% a year	IDC, Sep 2026
Latin America	About US$10 bn in 2026; Brazil US$4.2 bn	Mexico US$2.4 bn	IDC via EuQueroInvestir, Apr 2026
Middle East, Türkiye, Africa	US$7.2 bn in 2026; 37% a year	Fastest growth from the lowest base	IDC, May 2024

Regional forecasts come from different releases and a narrower definition than the worldwide total; they cannot be added up into shares.

© 2026 ETHORITY & STEN.AI LLC · Good AI Prism Study 2026/27 · ethority.ai/ai-prism-study

Where the private capital goes. Private AI investment reached US$285.9 billion in the United States in 2025, more than 23 times China's US$12.4 billion; the United Kingdom followed with US$5.9 billion. Organisational AI use rose to 88% of surveyed organisations, from 78% in 2024 (Stanford AI Index 2026, via Burges Salmon).

WHAT THIS SHOWS

Spending is rising fast and regular use is widely reported; an EBIT effect is reported by a minority of organisations.

WHAT IT MEANS FOR YOU

Budget alone does not buy results. The distance between trying AI and running it in a core process is where value is won or lost.

WHAT TO DECIDE NOW

Name the one process where AI in production would matter most to your results, and its accountable owner.

2 · ADOPTION AROUND THE WORLD

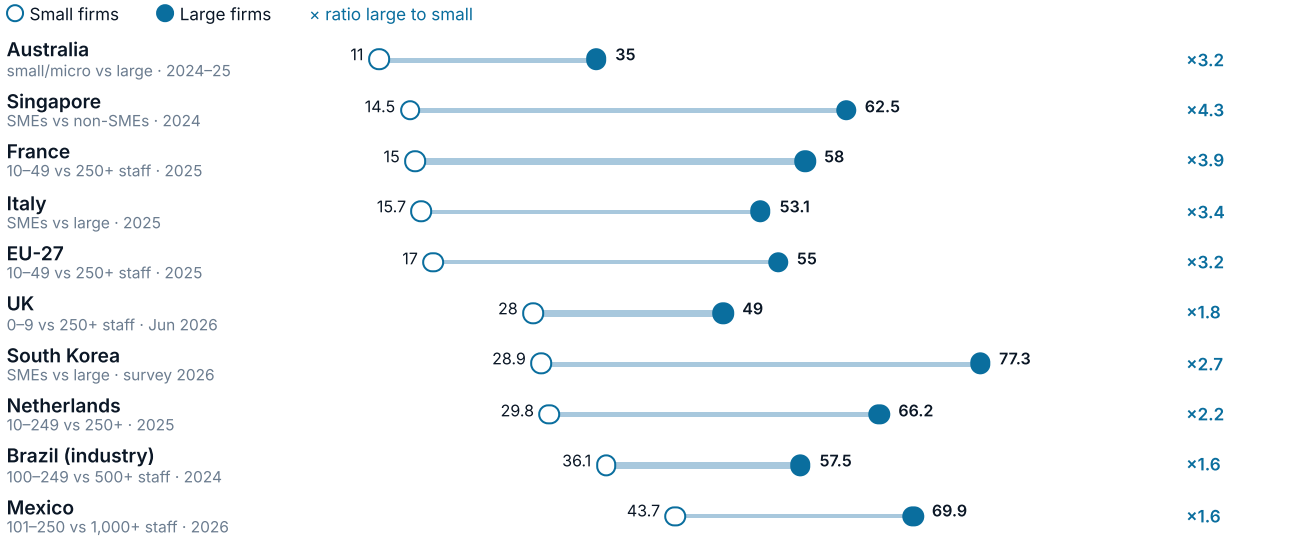
Two measures, and a size gap in every study

Firms and people are counted differently, and the two pictures do not always agree; within every national study we compared, large firms use AI more often than small ones.

EXHIBIT 5

In the selected studies, large firms adopt more often than small ones, by 1.6 to 4.3 times

Share of firms using AI, by size band, %; size bands and populations differ by study, so compare within a row, not across rows



Sources: ABS, IMDA, INSEE, Istat, Eurostat, ONS, Hankook Research for JoongAng, CBS, IBGE via Poder360, Banxico via Ecosistema Startup. Size bands differ by country, as labelled; rows are not a ranking.

© 2026 ETHORITY & STEN.AI LLC · Good AI Prism Study 2026/27 · ethority.ai/ai-prism-study

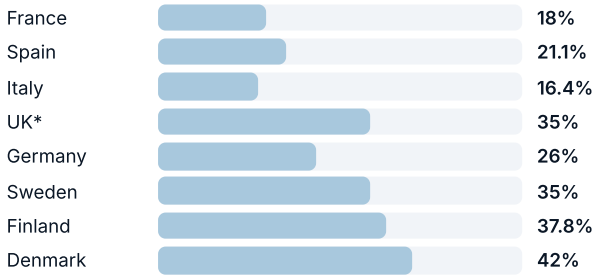
EXHIBIT 6

Two measures, two pictures: firm adoption and population use

Separate panels: different populations, definitions and periods; not a comparison of the same organisations

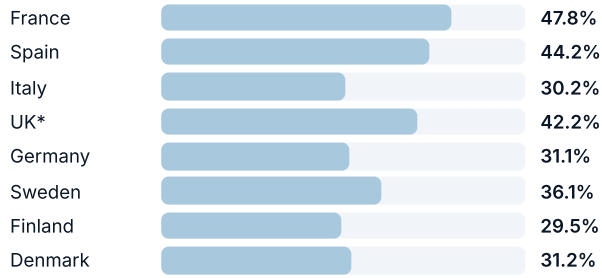
Firms with 10+ staff using AI, 2025

Official statistics (enterprise survey)



Working-age people using generative AI, Q1 2026

Vendor usage telemetry (Microsoft)



Sources: Eurostat and national statistics offices as in Exhibit 3; Microsoft. *UK firms: 10+ employees, June 2026 (ONS). The panels cannot show whether a firm's own staff use AI before the firm does.

© 2026 ETHORITY & STEN.AI LLC · Good AI Prism Study 2026/27 · ethority.ai/ai-prism-study

WHAT THIS SHOWS

In every national study we compared, large firms report AI use more often than small ones. Population use of generative AI and firm adoption are separate measures that point in different directions in several European economies.

WHAT IT MEANS FOR YOU

Small firms face a capability and cost gap that large firms have closed faster. Wide personal use means employees may meet AI at work whether or not the firm has a policy; the data cannot say how often.

WHAT TO DECIDE NOW

Small firms: set a short AI use policy and one sanctioned use case with clear data rules. Large firms: include small suppliers in your AI programmes.

3 · EUROPE

From one firm in twenty to two in five

Europe spans the widest range of adoption in the study: from 5% of firms in Romania to 42% in Denmark. The barriers are the same everywhere: expertise first, unclear law second. The EU AI Act gives the region a shared timetable.

ECONOMY	WHERE IT STANDS	MAIN PROBLEM	WHAT HELPS
EU average	20% of firms with 10+ staff use AI (2025), up 6.5 points; small firms 17%, large 55% (Eurostat)	Lack of expertise (71%) and unclear legal consequences (53%) among firms that considered AI	Practical skills programmes; plain-language AI Act guidance; use the Omnibus delay to build governance
United Kingdom	About 35% of firms with 10+ staff (June 2026), up from about 12% in 2023 (ONS)	71% of firms see "no identified need"; 60% lack skills (DSIT)	Use-case discovery for small firms; clarity from sector regulators
France	18% (2025), tripled since 2023; 58% of large vs 15% of small firms (INSEE)	71% of non-users see no use for AI; 54% lack expertise	Diagnostics and AI ambassadors under Osez l'IA, with a target of 80% SME adoption by 2030 (Ministère de l'Économie)
Italy	16.4% (2025), doubled in a year; SMEs 15.7% vs large 53.1% (Istat)	Skills (58.6%), unclear law (47.3%), data quality (45.2%)	Sector rules under the national AI law; SME skills programmes
Spain	21.1% of firms with 10+ staff and 13.4% of micro firms (2025) (INE)	Uneven spread: construction 11.4%; a national sanctions regime in preparation	A single enforcer and sandbox (AESIA)
Netherlands	SMEs 29.8%, large firms 66.2%, micro firms 13.8% (2025) (CBS)	Micro firms lag; several regulators share oversight	A clear map of who supervises what
Nordics	Denmark 42%, Finland 37.8%, Sweden 35%: the EU's top three (Eurostat)	Even leaders lack knowledge: 80% in Denmark (Danmarks Statistik)	Moving from use to measured productivity

Germany in depth (Mittelstand and large-enterprise needs, cost drivers, sovereignty) is in the Germany and US Focus edition and in the evidence annex.

REGIONAL CONCLUSION · EUROPE

OBSERVED CONSTRAINT	OUR INTERPRETATION	NEXT DECISION	2027 IMPLICATION
Lack of expertise (71%) and unclear legal consequences (53%) among EU firms that considered AI; adoption from 5% of firms in Romania to 42% in Denmark.	The European gap is capability and legal confidence, not access to tools. The Nordic figures show that organised adoption at two in five firms is reachable.	Pair skills support with a plain-language map of the AI Act duties that apply to your first use case.	Annex III duties for listed uses apply from 2 December 2027. Organisations that inventory and classify their AI in 2026 will meet them with less disruption.

4 · THE AMERICAS

The United States leads on spend; Latin America grows from a low base

US buyers ask "which tool, and does it pay?" Casual use is near-universal; integrated, profitable use is rare. Small firms need help choosing; enterprises need proof of profit, agent governance and cost control.

SEGMENT	STATUS	SOURCE
All US firms	19.8% use AI to produce goods or services; 37% at 250+ employees, under 20% at 1–4; no significant growth below 20 employees since Dec 2025	Census BTOS, May 2026
Small businesses	76% use AI, only 14% have it integrated into core operations (n=1,256)	Goldman Sachs 10KSB, Mar 2026
SMBs (QuickBooks base)	77% use AI regularly; 43% report higher revenue; only 12% pay for dedicated AI tools, but 86% of payers renew	Intuit AI Impact Report, May 2026
Solopreneurs	86% try an AI tool before paying a person; about 65% find checking AI output a drag (n=500)	FreshBooks / Wakefield, Sep 2026
Enterprises	Gen-AI spend US$37 bn in 2025 (3.2× 2024); 76% of use cases bought, not built; 27% of app spend product-led	Menlo Ventures, Dec 2025
Enterprises	2025 survey (1,993 respondents, 105 countries, fielded Jun–Jul 2025): 88% used AI, 39% saw any EBIT impact, mostly under 5%; about 6% high performers. August 2026 edition (1,719 respondents, 97 countries): 37%	McKinsey State of AI 2025
Large firms (US$1 bn+)	53% deploy agents; 35% cite AI cost management as a barrier; 26% have real-time cost visibility	KPMG AI Pulse Q2 2026

United States: needs, problems and responses

NEED	BUSINESS PROBLEM	WHAT HELPS	OUTCOME TO MEASURE
Choose the right tool	48% of small-business AI users find tool choice hard; 49% lack expertise (Goldman Sachs)	Task-based comparison with normalised pricing	Time to a confident choice; tools kept after 90 days
Trust the output and the data	Privacy and security top barrier (50%)	Data-use and hosting documentation for every tool in use	Data-handling terms checked for every tool in use
Prove profit, not pilots	Only 37% attribute some EBIT effect to AI (McKinsey, Aug 2026)	Business case with a decision gate	Cases that reach production; EBIT attributed
Govern agents and spend	Only 21% have mature agent governance (Deloitte); over 40% of agentic projects expected to be cancelled by 2027 (Gartner)	AI inventory, dependency map and governance	Inventory coverage; cost per outcome; cancelled-project rate
Meet new state rules	Colorado (SB 26-189) and California (CPPA) automated-decision duties from 1 January 2027 for covered uses (see section 4)	Readiness check: automated-decision notices, human review, records	Notices, opt-out route and records in place before the compliance date

REGIONAL CONCLUSION · UNITED STATES

OBSERVED CONSTRAINT	OUR INTERPRETATION	NEXT DECISION	2027 IMPLICATION
Integration and proof of profit: 14% of small businesses have AI in core operations; 37% of organisations attribute some EBIT effect to AI (McKinsey, Aug 2026); 26% of large firms see AI costs in real time.	The US market is a market of choice and return. Usage-based pricing makes cost control the next management task.	Compare tools for one job on total cost at your volume; make usage costs visible per use case before adding agents.	Colorado and California rules on automated decisions apply from 1 January 2027 for covered uses; map automated decisions now.

5 · ASIA-PACIFIC

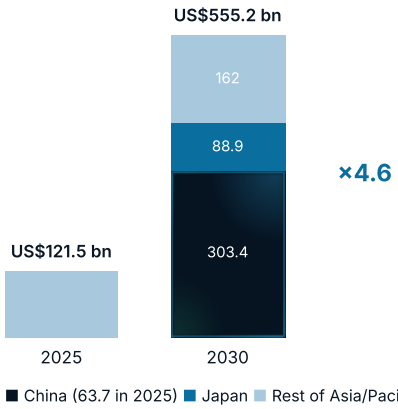
Fourfold growth forecast, and the widest size gap in our sample

Asia/Pacific AI spending is forecast to grow more than fourfold by 2030, from a base far larger than that of the Middle East and Africa. Firm-level data show the same size gap as in Europe, sharpest in Korean manufacturing.

EXHIBIT 13

Asia/Pacific AI spending grows 4.6-fold by 2030

US$ billion; China, Japan and the rest of the region



Source: IDC, Sep 2026 (forecast). Segments may not sum exactly due to rounding.

© 2026 ETHORITY & STEN.AI LLC · Good AI Prism Study 2026/27 · ethority.ai/ai-prism-study

35.5%
annual growth of Asia/Pacific AI spending, 2025–2030

16%
of large Japanese firms use AI company-wide

4.2%
of small and mid-sized Korean manufacturers use AI

42.8%
of China's population used generative AI by December 2025

Asia-Pacific by economy

ECONOMY	WHERE IT STANDS	MAIN PROBLEM	WHAT HELPS
China	602 million generative-AI users, 42.8% of the population (Dec 2025) (CNNIC); 98.1% of surveyed firms use generative AI (MIC Japan)	From broad use to measurable value; labelling and security compliance	Industry-specific models; label-ready content pipelines
Japan	86.4% of firms use generative AI for some task, but only 16% company-wide (MIC; Reuters/Nikkei)	Shallow use; nearly 30% of firms have no organisational effort behind it	Use policies (68.9% now have one); know-how transfer; use-case catalogues for SMEs
India	47% of large firms have several generative-AI use cases live; 10% are scaling (EY)	Pilot-to-production gap; data quality; data-protection duties from May 2027	Fast-to-deploy solutions; consent and data handling ready for the DPDP Rules
South Korea	77.3% of large firms vs 28.9% of SMEs use AI at work (Korea JoongAng Daily)	Only 4.2% of small and mid-sized manufacturers use AI; 80.7% have no AI talent (KCCI)	Supplier programmes run by large firms; compliance kits during the AI Basic Act grace period
Singapore	Non-SMEs 62.5% vs SMEs 14.5% (2024); 73.8% of workers use AI tools (IMDA)	SME gap; governance of AI agents	Agentic AI framework; AI Verify testing; MAS guidelines in finance
Australia	About 12% of all businesses, 35% of large and 11% of small firms (2024–25) (ABS)	Trust is the top SME barrier (about 65%); 54% see AI as not relevant (NAIC)	Six practices for AI adoption; worked use cases by industry
South-East Asia	65% of firms in a regional sample use some AI, 15% at an advanced level (UOB, bank survey)	Data and system readiness (47%), funding (47%), talent (39%)	SME grants; sandboxes; local data centres; Vietnam's law in force

REGIONAL CONCLUSION · ASIA-PACIFIC

OBSERVED CONSTRAINT
A sharp size gap (South Korea: SMEs 28.9% vs large 77.3%; 4.2% of small manufacturers) and shallow company-wide use (16% of large Japanese firms).

OUR INTERPRETATION
Scale is concentrated in large firms and in China; small suppliers risk falling behind the customers they serve.

NEXT DECISION
Large firms: run supplier AI programmes. All firms: adopt a written AI use policy, as 68.9% of Japanese firms now have.

2027 IMPLICATION
Grace periods end in 2027 (Vietnam 1 March; health, education and finance 1 September) and India's core data-protection duties apply from 13 May 2027.

Canada and Latin America

ECONOMY	WHERE IT STANDS	MAIN PROBLEM	WHAT HELPS
Canada	National AI Strategy "AI for All" announced June 2026	SME adoption and compute capacity	Privacy modernisation, SME programmes, sovereign compute (ISED)
Brazil	41.9% of industrial firms with 100+ employees used AI in 2024, up from 16.9% in 2022 (IBGE, PINTEC Semestral, Sep 2025)	Among firms not using advanced digital technologies, 74.3% cite high costs and 60.6% a lack of qualified staff (IBGE); legal uncertainty while the AI bill stalls	Risk-based law once passed; industrial programmes
Mexico	48.5% of firms with 100+ workers used at least one AI technology in June 2026, against 24.3% in September 2025 on a differently worded question (Banco de México, 2026)	Only 7% of larger firms report consolidated ROI; 46% of Mexican respondents struggle to understand AI cost structures (KPMG Global AI Pulse Q2 2026; firms with revenue above US$50 m)	Cost transparency per use case

Latin America is forecast to invest about US$10 bn in AI in 2026, with Brazil accounting for 42% of it (IDC). Brazil's AI bill has stalled in the Chamber of Deputies, so firms plan for a risk-based law without a date (CASRAI).

6 · MIDDLE EAST AND AFRICA

Highest personal use, thinnest firm data

The UAE has the highest share of people using generative AI in the world, and Saudi Arabia made the biggest jump in the rankings. For neither did we find official firm-level adoption figures. In Africa, use among formal firms is spreading quickly but remains shallow.

ECONOMY	WHERE IT STANDS	MAIN PROBLEM	WHAT HELPS
UAE	73.3% of working-age people use generative AI, the highest rate worldwide (Microsoft); 85% of CEOs say their culture enables AI (PwC)	Layered rules: federal, central-bank and free-zone	The new Federal Authority for AI and Data; enterprise ROI measurement
Saudi Arabia	Biggest rank gain in population use, to 25th worldwide (29.4% in Q1 2026); private AI investment US$2.03 bn in 2025	Binding data-protection law alongside non-binding AI guidance	Risk framework and the draft Global AI Hub Law
South Africa	67% of large enterprises had adopted generative AI in 2025, up from 45% in 2024 (World Wide Worx with Dell Technologies and Intel, n>100, July 2025; vendor-sponsored)	No binding framework yet	A phased, risk-based national AI policy
Nigeria · Kenya	In Nigeria, about 40% of formal firms with 5+ employees used AI in 2025, up from under 5% in 2022, with a similar rise reported for Kenya (World Bank, WDR 2026, figures as reported by Nairametrics)	Shallow use and dependence on foreign models; AI investment about 2% of revenue in Africa (PwC)	Mobile-first tools; depth of use, not only incidence

Middle East, Türkiye and Africa AI spending is forecast at US$7.2 bn in 2026, growing about 37% a year (IDC). The UAE's national strategy targets an economic contribution of AED 335 bn by 2031 (Morgan Lewis).

REGIONAL CONCLUSION · LATIN AMERICA · MIDDLE EAST AND AFRICA

OBSERVED CONSTRAINT

In Brazil, firms without advanced digital technologies cite cost (74%) and a lack of qualified staff (61%); 7% of larger firms report consolidated AI ROI (KPMG). In the Gulf and Africa, firm-level statistics are thin.

OUR INTERPRETATION

Growth comes from a low base and is cost-sensitive. High personal use in the Gulf does not, by itself, show enterprise depth.

NEXT DECISION

Ask which population and definition a figure covers; build cost transparency per use case before scaling.

2027 IMPLICATION

Brazil awaits its AI law; the UAE has a new federal authority; Saudi Arabia consults on a draft law. Plan for evolving, risk-based frameworks.

7 · SECTORS

Across the sectors we reviewed: which AI, how safely, and does it pay?

In the sectors we reviewed, the surveys describe organisations that are past "whether" and are now asking which AI, how safely, and whether it pays. Where an intended use is listed in Annex III of the AI Act, fixed dates apply; elsewhere the test is mainly return and risk. The response repeats with sector facts: an evidence-based choice, a readiness or audit step, and a measured outcome.

SECTOR	WHERE IT STANDS	BUSINESS PROBLEM	WHAT HELPS	OUTCOME TO MEASURE
Health & Life Sciences	81% of US physicians use AI, mostly documentation (AMA, Mar 2026); fastest-growing AI spender in Europe, 41% a year (IDC)	Documentation burden; validated safety; health-data rules	Choose documentation and triage tools on documented medical-device and data-handling facts	Clinician time saved; share of tools with documented clinical validation
Financial Services & Insurance	92% of EU banks deploy AI (EBA); banking is Europe's largest AI spender	Creditworthiness assessment and life and health insurance pricing are listed Annex III uses; DORA third-party risk	AI Act readiness for credit and insurance models; supplier-dependency map for third-party risk	Audit readiness by Dec 2027; fraud loss; cost to serve
Manufacturing & Industrial	88% of US mid-market manufacturers have integrated AI, 56% use agentic AI (RSM, 2026); Prism: 87 offerings researched in depth	Data quality behind 51% of failed pilots; unclear ROI behind 30%	Use-case portfolio (maintenance, inspection, copilots) with a data-readiness check and stop/go gates	Downtime, scrap rate, OEE
Retail & Consumer	About 14% of US retail firms use AI (Census); about 50% of consumers uneasy with autonomous AI purchases (Bain)	AI-driven discovery shifts traffic; trust in agentic commerce	Choose service, personalisation and pricing tools on documented facts; prepare for AI-driven discovery	Conversion, AI-referred traffic, forecast accuracy
Energy & Utilities	81% of North American utilities use AI (Itron); 34% of energy AI initiatives terminated vs 26% across industries (S&P 451)	High failure rate; safety components of critical infrastructure are listed in Annex III	Business case with a stop/go gate before scaling; governance for critical infrastructure	Outage minutes, forecast error, projects that reach ROI
Mobility & Logistics	Active AI use rose from 16% to 43% in a year; only 7% highly trust AI output (Logistics Management, Sep 2026)	Trust, document fraud, shadow AI (65% use AI with or without approval)	AI inventory to surface shadow AI; quality checks that build trust in AI output	Cost per shipment, on-time rate, fraud losses
Public Sector	64% exploring GenAI, 90% plan agentic AI (Capgemini); 60% of OECD cases still pilots	Only 36% of EU public bodies feel AI Act-ready; sovereignty	Sovereignty criteria, procurement-ready evidence, AI Act readiness review	Processing time, backlog, compliance
Education	60% of US teachers used AI; weekly users save 5.9 hours a week (Gallup-Walton); 19% of universities have a formal AI policy	Admission and assessment uses are listed in Annex III; missing policies; minors' data	A written AI policy, staff and student AI literacy, checks for high-risk assessment uses	Teacher hours saved; policy in place
Media & Entertainment	97% of news leaders see AI automation as important, 13% call results transformational; search traffic expected to fall 43% in three years (Reuters Institute)	Traffic loss to AI answers; licensing and IP	Rights and licence checks per tool; visibility in AI-generated answers	Cost per content unit, licensing revenue
Science & Research	84% of researchers use AI, up from 57% (Wiley); only 41% get adequate institutional support	Hallucinations (64% concerned); integrity; compute access	Open data and methods; documented model and data choices	Time to publication; reproducibility
Legal services	40% of professional-services organisations use GenAI, 15% agents (Thomson Reuters 2026); only 18% of firms track ROI	Clients question value; liability; shadow AI	Compare legal AI on documented accuracy; governance and ROI tracking	Hours saved, matter cost, ROI tracked

WHAT THIS SHOWS

Reported use is high in most sectors we reviewed, but measured results are thin and failure rates are material where reported.

WHAT IT MEANS FOR YOU

Your sector's adoption rate says little about your next decision. What matters is the intended use, the data it needs, the rules that apply and how you will measure it.

WHAT TO DECIDE NOW

Check the intended use of your first candidate process against Annex III and sector rules, and set the outcome measure before comparing tools.

10 · BUSINESS FUNCTIONS

Fourteen functions: high use, thin proof

Most functions show high reported use and thin measured proof. Where measured studies exist, gains are real but narrower than the marketing: +14% issues resolved per hour in customer support, and from 19% slower to 21% faster in software work.

FUNCTION	WHERE IT STANDS	BUSINESS PROBLEM	WHAT HELPS	OUTCOME TO MEASURE
Customer Service	91% of leaders under pressure to implement AI (Gartner); measured +14% resolutions per hour, +34% for novices (NBER)	Supplier noise, incomparable pricing, data-use risk	Compare tools on pricing unit and data use; pilot against resolution metrics	Cost per resolution, resolution rate, CSAT
Assistants & Copilots	UK government Copilot trial: 72% satisfied, no robust productivity evidence (The Register)	Licences bought, value unproven	Inventory licences; adoption and literacy programme with usage targets	Active use, hours saved, cost per active user
Knowledge & Search	Agents most used in IT and knowledge management (McKinsey)	Staff do not trust internal AI answers	Compare search and retrieval tools on hosting and answer accuracy	Answer accuracy, time to information
Documents, Workflow & Agents	No function above 10% scaled agents (McKinsey); over 40% of agentic projects expected to be cancelled by 2027 (Gartner)	"Agent washing", pilots that never ship	Business case with a stop/go gate; tightly scoped agent permissions	Cycle time, straight-through rate, cases in production
Sales & CRM	87% of sales organisations use AI (vendor survey, Salesforce)	Sellers spend 40% of time selling; disconnected systems	Check CRM and data integration before buying	Selling time, win rate, pipeline per rep
Marketing & Content	27% of CMOs report limited or no GenAI use in campaigns (Gartner)	Spend without proven payoff; brand and IP risk	Choose tools by job; check IP and licence terms	Content cost per asset, conversion
HR & Recruiting	Field experiment (about 70,000 applicants): AI interviews raised starts by 18% (Erasmus); EU high-risk duties from 2 Dec 2027 (Ogletree)	Recruitment and worker-management uses are listed in Annex III; bias; works-council co-determination	AI Act high-risk readiness; works-council involvement from the start	Time to hire, retention, audit readiness
Finance & Accounting	59% of finance functions use AI, flat year on year; 91% report low or moderate impact (Gartner)	Low impact despite use	Start with payables and close; measure against a baseline	Days to close, cost per invoice, DSO
Legal & Compliance	Hallucination rates of 17–33% in legal research tools (LLRX)	Accuracy and liability	Require source citations and evidence classes; human review of outputs	Verified-citation rate, review time
IT & Security Operations	Highest reported cost benefits (McKinsey); shadow AI behind breaches at 20% of firms (IBM, vendor, via Help Net Security)	Shadow AI and agent sprawl	AI inventory and dependency map; shadow-AI policy	Mean time to resolve, incidents, shadow-AI share
Software Development	84% use or plan to use AI, 46% distrust its accuracy (Stack Overflow); measured effects from −19% (METR) to +21% (Google RCT)	Speed claims versus quality and stability	Compare coding tools on data and licence terms; measure lead time and failure rate	Lead time, change failure rate
Data & Analytics	About 90% of data leaders not yet using the full AI capability set (Gartner, via vendor summary)	Data readiness, trust in numbers	Data-readiness and semantic-layer checks before scaling	Decision latency, forecast error
Design, Image & Video	86% of creators use generative AI (Adobe, vendor); IP law unsettled (Cleary Gottlieb)	Licence and IP risk	Commercial-use and training-data checks per tool	Cost per asset, rights incidents
Meetings, Voice & Language	33% of US workers have met with an AI notetaker; 65% of them say consent was inconsistent or missing (survey via KRDO)	Consent and recording law	Consent and hosting checks per tool	Meeting time saved, consent compliance

Pricing is shifting from seats to usage and outcomes: Gartner, cited by Deloitte, expects at least 40% of SaaS spend to be usage-, agent- or outcome-based by 2030. Comparing prices across units becomes a core buying skill.

WHAT THIS SHOWS

Measured gains exist, but they are specific to tasks and people, and some studies find no gain or a loss.

WHAT IT MEANS FOR YOU

A supplier's productivity figure is not your baseline. Your own measurement before and after a pilot is the number a board should rely on.

WHAT TO DECIDE NOW

For the function you start with, write down today's cost, volume, quality and cycle time before any pilot begins.

© 2026 ETHORITY & STEN.AI LLC · Put AI to Work: The Good AI Prism Global Market Study 2026/27 · **ethority.ai/ai-prism-study**

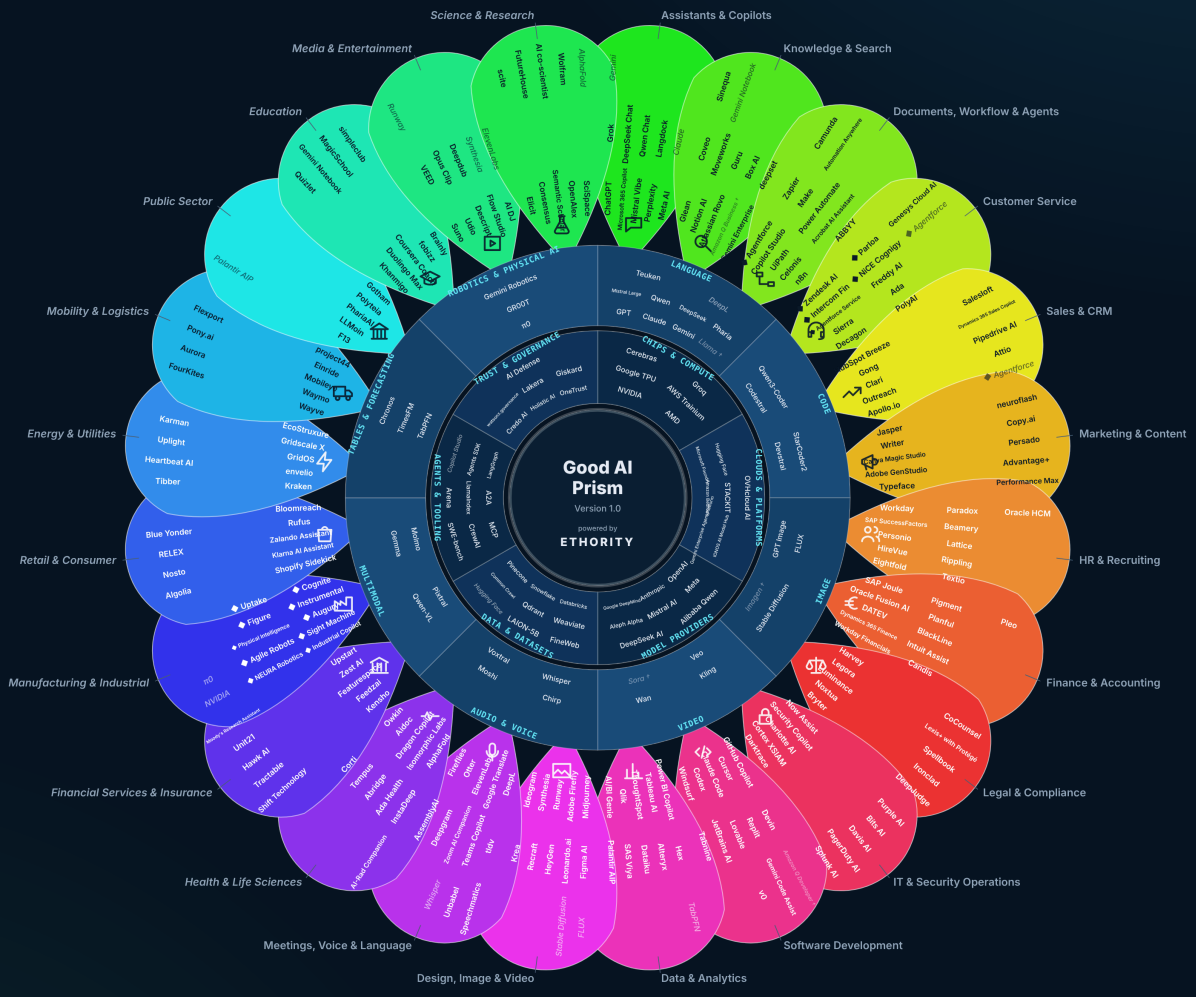
ETHORITY®

AI BUSINESS OPERATING SYSTEM

THE GOOD AI PRISM

Three rings. One AI market.

The Good AI Prism maps the AI market from the outside in. The outer ring shows where AI does work: 14 business functions and 10 industries. The middle ring shows the kind of data it handles, the inner ring the technology underneath. For this study we mapped 998 market entries, 961 of them AI offerings and 33 provider records (4 await classification), each recorded with a dated source from the provider's own pages.



961	24 · 8 · 6	194	11 · 14
AI offerings mapped from providers' own pages	business areas · data types · technology layers	depth-research records by research ID: Customer Service and Manufacturing	sectors and business functions analysed

Names on the wheel are an illustrative selection from the map. Colour and position show where AI does work; colour strength shows research depth. An entry can appear in more than one area; the map counts 1,037 placements for 998 entries.

TAXONOMY AND UNITS

What is counted, and what is not

A market map is only as useful as its units. Company, product, configuration, placement and research record are different things; this edition keeps them apart and never adds them up.

UNIT	DEFINITION	THIS EDITION	NOT THE SAME AS
Market entry	A named item on the map with a dated source from the provider's own pages	998	A verified product or a customer
AI offering	An entry that is an AI product or service (incl. products named after their provider)	961	A distinct commercial product; overlaps are possible
Provider record	An entry that describes a company rather than an offering	33 (+4 awaiting classification)	A ranking of providers
Placement	A position of an entry on the wheel; one entry can sit in several areas	1,037	A count of offerings
Depth-research record	Decision fields assembled for a buyer task (price unit, hosting, training use, integrations, limits)	194 (Customer Service 107, Manufacturing 87)	Measured performance
Decision value	One sourced field value in a research record	775	An independent assurance
Relationship	A recorded dependency between offerings (e.g. built on a model or platform)	216	A partnership claim
Configuration	Product × plan × region × date: the unit a buyer actually signs	Not counted in Wave 1	An offering
Extension research set	A sector map assembled for a theme study (e.g. care, pharma)	Separate; not admitted to Wave 1	Additive market totals

The taxonomy of the wheel

Outer ring · 24 business areas

14 functions: Customer Service; Assistants & Copilots; Knowledge & Search; Documents, Workflow & Agents; Sales & CRM; Marketing & Content; HR & Recruiting; Finance & Accounting; Legal & Compliance; IT & Security Operations; Software Development; Data & Analytics; Design, Image & Video; Meetings, Voice & Language. 10 industries: Health & Life Sciences; Financial Services & Insurance; Manufacturing & Industrial; Retail & Consumer; Energy & Utilities; Mobility & Logistics; Public Sector; Education; Media & Entertainment; Science & Research.

Middle ring · 8 data types

Language, voice, audio, image, video, tables, code and multimodal: the kind of data an offering handles.

Inner ring · 6 technology layers

Models, chips and compute, clouds, data, agents and trust: the technology underneath.

Sectors in this study · 11

The 10 industries plus legal services, which the wheel maps as a function. HR is a function; care and pharma are sub-sectors of Health & Life Sciences.

MARKET MAP · WHAT THE MARKET DOCUMENTS

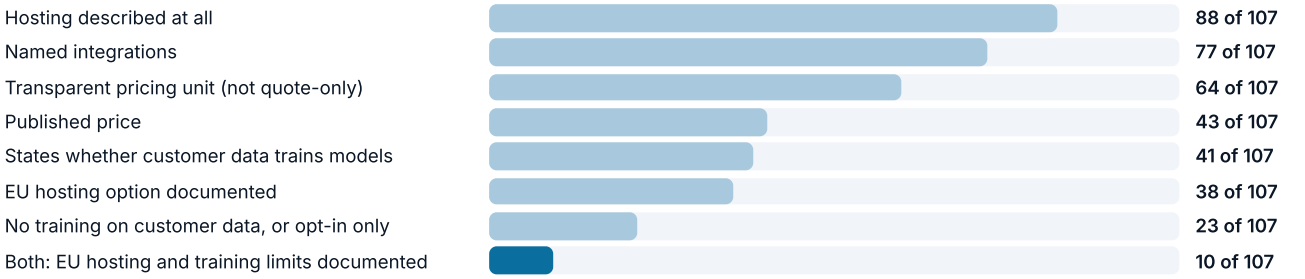
The facts a careful buyer needs are mostly not public

Of 107 customer-service AI offerings reviewed, we found public documentation for both criteria in 10: an EU hosting option, and either a statement that customer data is not used to train the provider's models (8 offerings) or that it is used only with opt-in consent (2). The two kinds of disclosure are not identical guarantees. The count was recomputed by offering ID from the frozen 107-row fact matrix of the research edition.

EXHIBIT 14

In 10 of 107 offerings we found public documentation of both facts

Customer-service AI offerings whose own public pages document each fact; provider information reviewed by ETHORITY on 19–20 October 2026; Wave 1 of ETHORITY's longitudinal research programme



Missing documentation does not mean a capability is missing. 'Training limits': customer data not used to train the provider's models (8 offerings) or used only with opt-in consent (2).

© 2026 ETHORITY & STEN.AI LLC · Good AI Prism Study 2026/27 · ethority.ai/ai-prism-study

WHAT THIS SHOWS

For a European buyer with data-protection duties, two facts decide the shortlist. Ten offerings document both; for 97, the public pages leave at least one open.

WHAT IT MEANS FOR YOU

The shortlist is decided less by features than by facts most providers do not publish. Not documented is not the same as not available: many of the 97 may qualify, but the buyer has to ask.

WHAT TO DECIDE NOW

Write both conditions into the request for proposal before any demo. Ask for the hosting region and the training-use clause in writing; treat an unanswered question as an open risk.

The same pattern runs across the market

- **Marketing, not documentation.** Four in five offerings we added in 2026 are described only on marketing pages, not in technical documentation.
- **Price is hard to compare.** Customer-service offerings bill in at least seven different units: per resolution, per seat, per conversation, per minute, credits, plan tier and quote only. Without a common unit, prices cannot be compared.
- **Dependencies are hidden.** Verified relations between offerings show applications resting on a few hubs: Zendesk, Salesforce and the Model Context Protocol on the application side; OpenAI, Anthropic, Amazon Bedrock and Hugging Face underneath. One model or platform change can move many products at once.
- **Identity moves fast.** 34 renames and 26 ownership changes surfaced within a few weeks of research, for example Acrolinx to Markup AI and LivePerson to SoundHound AI.

BOUNDARY CONDITIONS

Barriers, dated rules and the limits of this evidence

Three conditions bound every AI decision in this study: the barriers firms name, the rules that apply on fixed dates, and what this edition could and could not observe.

EXHIBIT 7

Lack of expertise tops the list in most economies we compared

Share of firms naming each barrier; bases differ by country

ECONOMY AND BASE	LACK OF EXPERTISE	LEGAL UNCERTAINTY	COST	SOURCE
EU-27 firms that considered AI	70.9%	52.5%	–	Eurostat 2025
Germany non-users	85%	66%	63%	Bitkom 2026
Denmark considered, not using	80%	–	–	Danmarks Statistik
Italy non-users	58.6%	47.3%	43%	Istat 2025
France non-users	54%	38%	32%	INSEE 2026
United Kingdom firms with 5+ staff	60%	28%	23%	DSIT 2026
Brazil industry non-users, digital tech	60.6%	–	**74.3%**	IBGE via Poder360
South Korea manufacturers	80.7%	–	73.6%	KCCI 2025

A dash means the barrier was not asked in comparable form. Surveys are not pooled. Brazilian industry is the exception: cost is named more often than expertise.
Sources: Eurostat, Bitkom, Danmarks Statistik, Istat, INSEE, DSIT, IBGE via Poder360, KCCI via JoongAng.

© 2026 ETHORITY & STEN.AI LLC · Good AI Prism Study 2026/27 · ethority.ai/ai-prism-study

Dated obligations a multinational plans against

In force	South Korea AI Basic Act (22 Jan 2026); Vietnam Law on AI (1 Mar 2026; grace to 1 Mar 2027); China content labelling (1 Sep 2025)
2 Dec 2026	EU: machine-readable marking for generative AI already on the market
1 Jan 2027	Colorado and California: duties for covered automated decisions
13 May 2027	India: core data-protection duties for businesses
2 Dec 2027	EU: high-risk duties for Annex III uses (employment, credit, insurance, education, public services)
2 Aug 2028	EU: high-risk AI in regulated products (Annex I), incl. machinery

Application dates from EUR-Lex (Reg. (EU) 2026/1744) and the national sources listed in the evidence annex; not legal advice.

What this edition does not show

- No official firm-level adoption statistic was found for the UAE, Saudi Arabia, India or China.
- Depth research covers Customer Service and Manufacturing only; other areas are mapped, not profiled.
- Provider pages were read on 19–20 Oct 2026; later changes are recorded in Wave 2.
- A missing published fact is an open question, not an absent capability.
- Regional spending forecasts come from different releases and are not added up.
- Telecoms, construction, professional services and hospitality have no industry petal yet.

WORKED CASE · ILLUSTRATIVE

CASE 1 · CUSTOMER SERVICE · BASELINE AND CHOICE

Choosing a customer-service AI on facts, not demos

A European service organisation with about 10,000 customer cases a month wants AI to resolve routine requests. Its data-protection officer requires EU hosting and no training on customer data. The market offers 107 profiled options.

Baseline

Volume	10,000 cases per month (illustrative)
Today's measure	Cost per resolved case, first-contact resolution, customer satisfaction, cycle time
Owner	Head of customer service (accountable); quality lead accepts outcomes
Constraint	EU hosting and limits on training with customer data, in writing

Options and deciding facts

STEP	EVIDENCE	RESULT
Screen the market	Exhibit 14: 10 of 107 offerings document both EU hosting and training limits	Shortlist of documented options; 97 others asked in writing
Normalise price	At least seven billing units: per resolution, seat, conversation, minute, credit, tier, quote	One cost per accepted outcome at the organisation's own volume
Check dependencies	Applications rest on a few hubs and model providers	Fallback named for each critical dependency
Agree the decision rule	Measured gains: +14% resolutions per hour, +34% for novices (NBER)	Threshold and stop condition written before any demo

> Decision: run one 8-week pilot with the best-documented option; keep the second as a tested alternative.

Evidence used in this case

Exhibit 14 (p. 17) · Billing units (p. 17) · NBER customer-support study (p. 14) · Cost per accepted outcome (p. 37)

Illustrative case built from the evidence in this study. Values marked "illustrative" show the method; they are not a client result.

WORKED CASE · ILLUSTRATIVE

CASE 1 · CUSTOMER SERVICE · ACCEPTANCE AND LEARNING

Choosing a customer-service AI on facts, not demos

Acceptance conditions

MEASURE	THRESHOLD (ILLUSTRATIVE)	WHO ACCEPTS
Accepted resolutions	≥ 35% of attempted cases resolved to the customer's acceptance	Quality lead
Quality floor	Complaint rate not above baseline; review sample of 5% of AI-resolved cases	Quality lead
Cost per accepted outcome	Below today's cost per resolved case on the same monthly boundary	Finance partner
Data terms	Hosting region and training clause confirmed in the contract	Data-protection officer

Cost boundary (illustrative, per month): usage €6,000; amortised integration €1,500; monitoring, change and training €500; human review €2,000; failed or escalated cases €1,000. With 40% accepted: €11,000 ÷ 4,000 = €2.75 per accepted outcome.

Possible outcomes

Proceed

Threshold met on accepted outcomes and quality; scale to further case types.

Redesign

Quality met, cost not: change billing unit, scope or routing.

Stop

Quality floor breached or terms not confirmed in writing.

Record for the next wave: which documentation was missing, what the provider confirmed in writing, and the measured cost per accepted outcome.

Decision record (summary)

Decision	Pilot option A for 8 weeks; option B kept as tested alternative
Owner	Head of customer service
Evidence	Written hosting and training terms; baseline of 3 months
Open risks	Billing unit changes; model change by the provider
Review date	End of week 8

ETHORITY route: Good AI Prism research and Buyer Intelligence for the shortlist; an Opportunity Case and Decision Record for threshold, cost boundary and owner.

Illustrative case built from the evidence in this study. Values marked "illustrative" show the method; they are not a client result.

WORKED CASE · ILLUSTRATIVE

CASE 2 · MANUFACTURING · BASELINE AND CHOICE

Quality support on the line: start with data, not with a model

A mid-sized component manufacturer wants AI-supported visual inspection to cut scrap without letting more defects reach customers. In manufacturing, data quality is behind 51% of failed AI pilots and unclear ROI behind 30% (RSM, 2026).

Baseline

Line	One machining line, two shifts (illustrative)
Today's measure	Scrap rate, defect escape rate to customers, false rejects, inspection hours
Owner	Plant manager (accountable); quality manager accepts outcomes
Constraint	Images stay on site; the supplier may not reuse them without consent

Options and deciding facts

OPTION	DECIDING FACTS TO OBTAIN	RISK TO TEST
Machine-vision system with AI	Edge deployment; retraining for new variants; data ownership of images	False rejects on new parts
AI module from the line-equipment supplier	Integration with the existing controller; update and support terms	Lock-in to one equipment roadmap
Quality copilot on existing process data	Data readiness of sensor and MES records; who labels defects	Recommendations nobody acts on

> Decision: a two-month data-readiness check first; only then a pilot of one option on one product family. The Prism profiles 87 manufacturing offerings in depth to build the shortlist.

Evidence used in this case

Manufacturing row (p. 13) · 87 depth records (p. 16) · Machinery Regulation (p. 33) · Robotics indicators (p. 33)

Illustrative case built from the evidence in this study. Values marked "illustrative" show the method; they are not a client result.

WORKED CASE · ILLUSTRATIVE

CASE 2 · MANUFACTURING · ACCEPTANCE AND LEARNING

Quality support on the line: start with data, not with a model

Acceptance conditions

MEASURE	THRESHOLD (ILLUSTRATIVE)	WHO ACCEPTS
Scrap rate	At least 15% below baseline over 6 weeks	Quality manager
Escape rate	Not above baseline (no extra defects reach customers)	Quality manager
False rejects	Within the agreed band; reviewed by an inspector	Shift lead
Safety and rules	If the AI acts as a safety function: EU Machinery Regulation from 20 Jan 2027 and AI Act Annex I from 2 Aug 2028 checked	Plant manager

Cost boundary: system or licence; cameras and edge hardware; integration with the controller; labelling of defect images; inspector review time; retraining for new variants; downtime during installation.

Possible outcomes

Proceed

Scrap down, escapes not up: extend to a second product family.

Redesign

Escapes stable but false rejects too high: retrain or narrow the scope.

Stop

Data not ready: fix data collection first; the model is not the bottleneck.

Record: labelled-image volume needed, retraining effort per new variant, and inspector time per shift before and after.

Decision record (summary)

Decision	Data-readiness check, then pilot on one product family
Owner	Plant manager
Evidence	Labelled image set; baseline scrap and escape rates
Open risks	New variants; image rights; safety classification
Review date	After the readiness check

ETHORITY route: a Check and Opportunity Case to test data readiness and economics; engineering and integration as a defined project scope.

Illustrative case built from the evidence in this study. Values marked “illustrative” show the method; they are not a client result.

WORKED CASE · ILLUSTRATIVE

CASE 3 · CROSS-BORDER ENTERPRISE · BASELINE AND CHOICE

One AI inventory for every jurisdiction

A group with operations in the EU, the United States, India and South Korea does not know how many AI systems it runs, including AI features inside software it already pays for. Its earliest obligations already apply.

Baseline

Scope	Four regions, about 40 business units (illustrative)
Today's measure	Share of AI systems with a named owner, risk class and deployment region; known shadow-AI use
Owner	Chief information officer (accountable); compliance accepts the register
Constraint	Earliest deadline per market: Korea in force; US states 1 Jan 2027; India 13 May 2027; EU 2 Dec 2027

Options and deciding facts

OPTION	WHAT IT COVERS	WHAT IT MISSES
Survey of business units	Fast, cheap first picture	Shadow AI and AI features in existing software
Discovery from contracts, SSO and network data	Paid and sanctioned tools, embedded features	Purpose and risk class
Combined register with owners	Owner, purpose, data, region, risk class, earliest deadline	Needs a refresh cycle to stay true

> Decision: combined register; each system gets an owner, a purpose, a deployment region and the earliest deadline that applies. Large German firms show the gap: 98% have an AI strategy, 37% clear governance (KPMG, 2026).

Evidence used in this case

Dated obligations (p. 18) · Regulation table (annex) · Germany: strategy vs governance (annex) · Architecture (p. 26)

Illustrative case built from the evidence in this study. Values marked "illustrative" show the method; they are not a client result.

WORKED CASE · ILLUSTRATIVE

CASE 3 · CROSS-BORDER ENTERPRISE · ACCEPTANCE AND LEARNING

One AI inventory for every jurisdiction

Acceptance conditions

MEASURE	THRESHOLD (ILLUSTRATIVE)	WHO ACCEPTS
Coverage	≥ 95% of systems found by discovery are in the register	Compliance
Ownership	Every system has a named owner and a purpose	CIO
Classification	Risk class and earliest deadline recorded per jurisdiction	Legal
Refresh	Quarterly update; new purchases enter before go-live	Procurement

Cost boundary: discovery tooling; owner interviews; legal classification; register maintenance; training for owners. The register is valued by avoided duplicate spend and deadline readiness, not by hours saved.

Possible outcomes

Proceed

Register complete: prioritise the systems with the earliest deadlines.

Redesign

Coverage low: add contract and network discovery.

Stop

No owners accept responsibility: escalate to the executive sponsor before continuing.

Record: share of AI found only by discovery, duplicate tools retired, and deadlines met on time.

Decision record (summary)

Decision	Combined register with owners and earliest deadlines
Owner	Chief information officer
Evidence	Discovery results; owner confirmations; legal classification
Open risks	Shadow AI outside discovery; owners not accepting
Review date	Quarterly

ETHORITY route: AI Estate to hold inventory, owners and dependencies (demonstration and early access); governance work as a scoped engagement.

Illustrative case built from the evidence in this study. Values marked “illustrative” show the method; they are not a client result.

DECISION JOURNEY

The decisions that will shape 2027

The evidence turns into an agenda for practical progress: establish the baseline, run a supported pilot, learn from the result and scale where the outcome justifies further investment.

EXHIBIT 17

From a first AI decision to a steered AI estate

Three phases and two gates

Now

BEFORE THE PILOT

- Choose one process that matters to results
- Name the person accountable for it
- Measure today: volume, cost, quality, time
- Write down, before any demo, the result that justifies a pilot and the one that stops it

Next 90 days

FROM YOUR START DATE

- Shortlist three credible options for that task
- Get the deciding facts in writing: billing unit, data location, training use, integration
- Run one pilot under a written charter
- Count accepted outcomes, not activity

2027

SCALE AND STEER

- Scale only what met its threshold
- Track cost per accepted outcome over time
- Keep one accountable owner per AI system
- Review the whole AI estate at least once a year

◆ **Gate 1: baseline agreed.** No pilot starts without a measured starting point and a named owner.

◆ **Gate 2: pilot meets its threshold.** No scale-up without the agreed result on accepted outcomes.

© 2026 ETHORITY & STEN.AI LLC · Good AI Prism Study 2026/27 · ethority.ai/ai-prism-study

EXHIBIT 18

Pilot, ask for more evidence, redesign or stop

Five questions, in order; each has a defined next step

1	Is the process important enough to change?	No → choose another process.
2	Do we have a measured baseline and an owner?	No → measure first (Gate 1).
3	Do the deciding facts exist in writing for at least two options?	No → request evidence before any demo.
4	Have the baseline checks been done, and is the intended use high-risk?	Every pilot: privacy, security, workplace, contract and sector rules. If the intended use is high-risk, add the additional readiness controls first.
5	Did the pilot meet its threshold on accepted outcomes?	No → redesign or stop. Yes → scale and steer (Gate 2).

© 2026 ETHORITY & STEN.AI LLC · Good AI Prism Study 2026/27 · ethority.ai/ai-prism-study

The cost of the decision is part of the decision: licences or usage, integration, human review and rework, divided by the outcomes the business accepts (see pages 28 and 37). Each phase ends with a decision the reader can name: proceed, request more evidence, redesign or stop. Value is judged against the objective that was set: financial return, service quality, rights protection or public benefit, over a review horizon that fits the risk.

ARCHITECTURE AND DEPENDENCIES

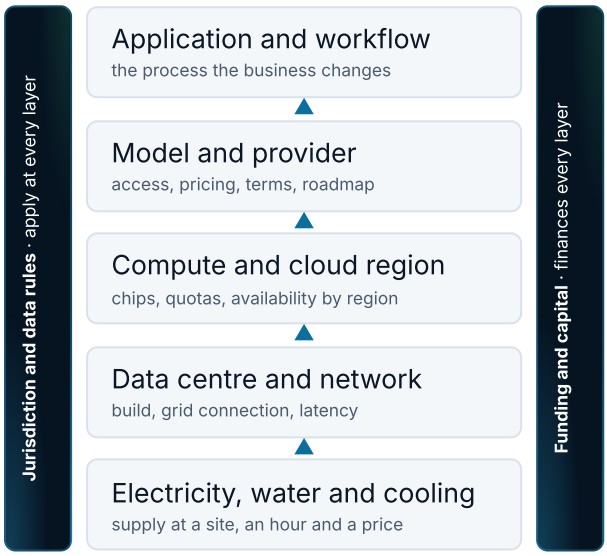
AI runs on physical and financial foundations

AI looks like software to the buyer, but it depends on chips, data centres, grids, water, capital and law. Hyperscalers are expected to spend more than US$700 billion in capital expenditure in 2026, mostly on new data centres, while grids, connection queues, permitting and electrical equipment set the pace (IEA).

EXHIBIT 15

What an AI application depends on

Each layer rests on the ones below; law and capital apply throughout



A dependency is not a forecast of failure. It shows where a change outside the organisation can reach a business process, and where to ask for evidence.

© 2026 ETHORITY & STEN.AI LLC · Good AI Prism Study 2026/27 · ethority.ai/ai-prism-study

EXHIBIT 16

Electricity used by data centres worldwide

All data centres, not AI alone; terawatt-hours per year



About 3% of global electricity demand in 2030. Use by AI-focused data centres rose 50% in 2025 alone.

Source: IEA, Key Questions on Energy and AI, 2026. Projection, base case.

© 2026 ETHORITY & STEN.AI LLC · Good AI Prism Study 2026/27 · ethority.ai/ai-prism-study

FINANCIAL CONDITIONS

The IMF's October 2026 Global Financial Stability Report warns that weaker or less durable investment and profitability in AI could trigger a correction in AI-linked assets, amplified by rising leverage and concentrated exposures. That is a conditional risk, not a forecast; for buyers it is a reason to check supplier continuity and exit terms.

From announcement to usable capacity

STATUS	WHAT IT SHOWS	WHAT TO ASK
Announced	Intent and headline value	Is there a site, a customer and a date?
Financed	Capital committed	By whom, and on what conditions?
Permitted	Planning and environmental approval	Which approvals remain outstanding?
Connection agreed	Grid connection contract signed	What capacity, from which date?
Energised	Power physically available at the site	Since when, and at what load?
Commissioned	Capacity operating	IT load, or total power?
Customer-available	Capacity a buyer can actually use	In which region, with what quota and SLA?

Announced amounts are not operating infrastructure, and capital expenditure is not spending on usable AI or customer return; the two must not be added to AI software spend. In Texas, Virginia, Finland and the UK, queued data-centre projects already exceed today's peak demand, and queues include speculative and duplicate requests. Enough electricity worldwide does not guarantee power at a given site, hour or price.

WHAT THIS SHOWS

AI capacity is being built at record pace, but physical and financial constraints decide where and when it becomes available.

WHAT IT MEANS FOR YOU

For most buyers the risk is not a power cut. It is capacity, price or terms in their region changing during a contract.

WHAT TO DECIDE NOW

For each critical AI use, record the region, the provider, the fallback and the contract terms for capacity and exit.

ABOUT THIS STUDY · EVIDENCE

Every number has a source

The figures describe the frozen research edition of 20 October 2026. Research depth and publication clearance are different dimensions: neither a populated profile nor a source check establishes operational performance.

1,037
placements on the Prism wheel

998
market entries, including 961 AI offerings, each with a dated provider source

194
depth-research records by unique research ID, including 12 earlier profiles

775
sourced decision values: pricing unit, price, hosting, integrations, prerequisites, limits

216
relationships between offerings recorded in this frozen edition

103
decision values re-checked in a separate research pass: 97 confirmed, 6 minor wording issues; not independent assurance

How to read the evidence

Each step needs recorded evidence. Wave 1 reaches step 4 for 194 records and step 5 for the Customer Service finding, whose pages ETHORITY read and reviewed on 19–20 October 2026; steps 6 and 7 are not claimed.

1	Observed	A source row with locator and date	Every entry
2	Identified	Name, provider and entity type recorded	998 entries
3	Inspected	Fields hold source excerpts and evidence classes	Research profiles
4	Research profile	Decision fields assembled for a buyer task; coverage can be incomplete and terms change	194 records
5	Human reviewed	ETHORITY reviews the pages behind the finding	Wave 1
6	Operating evidence	Results observed in a defined real-world context	Not claimed
7	Verification and assurance	Separate independent activities	Not claimed

Where the claims come from

Sources behind the 701 placements added in this edition. These counts describe sources, not independent validation.

Provider marketing pages	564 of 701
Provider documentation	137 of 701

How facts are kept

- Each fact keeps its page, excerpt and date, and the product, plan and region it applies to.
- A separate research pass re-opens the cited page for a random sample.
- Providers and readers can submit corrections at ethority.ai/ai-prism-study; confirmed corrections are dated in the next edition.

Unknown is a research question

A missing published fact does not mean a capability is absent. Ask the provider, record the answer, update the edition.

Colour shows depth

Position and colour on the wheel show where AI does the work and how deeply it has been researched.

Market and client data stay apart

Market research stays separate from an organisation's private AI estate and implementation evidence.

VALUE CONVERSION

Full cost and accepted outcomes

AI value is not the hours a supplier says it saves. It is the outcome the business accepts, at a full cost that includes integration, review and rework, compared with today on the same boundary.

From activity to attributable value

01 **Activity** — Prompts, sessions, licences used

02 **Time released** — Minutes no longer spent on a task

03 **Capacity used** — Released time put into service, quality or growth

04 **Accepted outcome** — A result the business accepts: resolved case, approved invoice

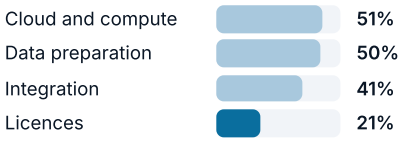
05 **Attributable value** — Change against baseline, net of full cost

COST ELEMENT	WHAT TO INCLUDE	OFTEN MISSED
Usage or licence	Seats, tokens, resolutions, credits at your own volume	Overage and price changes during the term
Integration	Connectors, data preparation, testing; amortised	Data preparation (50% of German firms name it a cost driver)
Infrastructure	Cloud, compute, storage in the chosen region	Cloud and compute (51% name it)
Human review	Time to check outputs and handle exceptions	Review effort in high-volume uses
Failure handling	Rework, escalations, complaints	Counting rework twice
Change and training	Literacy, process change, monitoring	Time of managers and staff

EXHIBIT 12

The licence is the smallest cost

Main AI cost drivers named by German firms



Source: Bitkom, Sep 2026

© 2026 ETHORITY & STEN.AI LLC · Good AI Prism Study 2026/27 · ethority.ai/ai-prism-study

WHAT THIS SHOWS

Licences are the smallest cost driver named by German firms; cloud, data preparation and integration cost more.

WHAT IT MEANS FOR YOU

A business case built on licence price and hours saved overstates value.

WHAT TO DECIDE NOW

Agree the cost boundary and the accepted outcome before the pilot; use the worksheet on page 37.

OPERATING MODEL

Workforce and operating model

AI changes tasks before it changes the organisation chart. Every system needs named people for the result, the review and the challenge route, and a decision about where released capacity goes.

ROLE	ACCOUNTABLE FOR	IN THE PILOT	IN OPERATION
Executive sponsor	Objective, budget, stop decision	Approves charter and gates	Reviews value and the AI estate yearly
Process owner	The result of the process	Sets baseline and threshold	Owns cost per accepted outcome
Outcome acceptor	What counts as an accepted result	Defines acceptance and samples	Signs off quality monthly
Oversight owner	Outputs that affect people	Reviews consequential outputs	Runs the challenge route
Data and system owner	Data rights, access, integration	Confirms hosting and training terms	Manages changes and incidents
Employee representatives	Participation where law or agreement requires	Consulted before start	Review effects on roles

Capability to build

- AI literacy for everyone who uses or reviews the system
- Review skills: when to trust, when to escalate
- Measurement: baseline, accepted outcomes, full cost
- Supplier management: terms, changes, exit

Evidence: 66% of German firms rate staff AI skills as low (Bitkom); 81% of Korean manufacturers have no AI specialists (KCCI); recruitment and worker-management uses are EU high-risk from 2 Dec 2027; Italy requires employers to inform staff when AI is used at work.

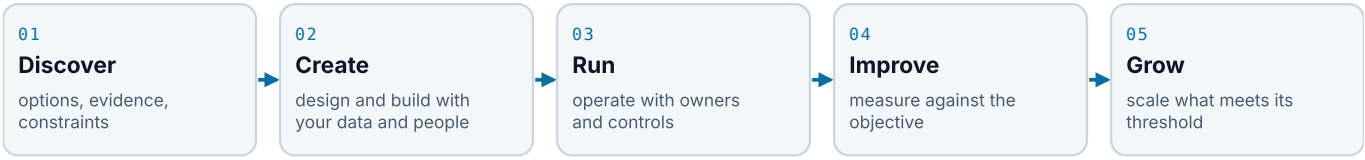
HOW ETHORITY HELPS

From this study to an accountable decision

Start with a decision that matters. The Good AI Prism helps you examine market options and their documented evidence. ETHORITY then connects the shortlist to your objectives, systems, people and economics: defining a pilot, building the capability, establishing responsibility and measuring accepted outcomes.

YOUR QUESTION	ETHORITY CAPABILITY	WHAT YOU RECEIVE
What exists and fits?	Good AI Prism, Registry, Navigator; Buyer Intelligence	Dated profiles, comparison and a reasoned shortlist
What should we commit to?	Check, Opportunity Case, Decision Record	Baseline, options, economics, owner, acceptance and stop rules
Where is the decision kept?	Mission / My ETHORITY Workspace	Objective, evidence, assumptions, decision and review date together
How do we coordinate systems?	AI Estate and governance work	Inventory, owners, dependencies and controls
How do we build and operate?	Transformation, engineering, governed workflows and agents	Tested implementation, human roles, operating support, measurement
What happens when conditions change?	Evidence Watch / TrustOps, research and review	Dated change, affected claims and decisions, accountable review
How could parties exchange trust?	Trust Profile, Passport, Handshake, Marketplace	Evidence scope and future participation

The ecosystem cycle



Delivery route for each use: decide → pilot → build → run → scale. Each capability is described at its actual delivery stage on the next page.

CAPABILITY STATUS

What is available today, and what proof stands behind it

A research study should not blur the line between a product that runs and a design that is being built. This page states each capability at its actual stage. Green marks what is available or recommended now; it does not mean independently verified.

01 **Published research** — read and cite today
→ 02 **Offered** — subscription or scoped engagement
→ 03 **Preview** — working, limited scope
→ 04 **Demonstration** — shown, not in operation
→ 05 **Design · roadmap** — planned or simulated

CAPABILITY	STAGE TODAY	WHAT STANDS BEHIND IT	BEFORE A STRONGER CLAIM
Good AI Prism, Registry, Navigator	PUBLISHED RESEARCH	This study; dated research edition of 20 Oct 2026; public map at release	Wave 2 comparison and correction log
Buyer Intelligence	RESEARCH SUBSCRIPTION	Research profiles and comparisons from the same evidence base	Scope and terms in the current offer sheet
Check, Opportunity Case, Decision Record	SCOPED ENGAGEMENT	Defined deliverables; acceptance by the client	Client-accepted results recorded per case
Mission / My ETHORITY Workspace	PREVIEW · DEMONSTRATION	Working preview of the decision record	Tenant isolation and persistence proven in operation
AI Estate	DEMONSTRATION · EARLY ACCESS	Inventory, owners and dependencies shown in a demonstration	Isolation, roles, security and support in place
Transformation and engineering	DEFINED PROJECT SCOPE	Project scope, roles and acceptance per engagement	No universal deployed platform claim
Evidence Watch / TrustOps	CONFIGURED DESIGN · SCOPED SERVICE	Dated review of changes within a contracted scope	Measured operating proof before any always-on claim
Trust Profile, Passport, Handshake, Marketplace	DESIGN · SIMULATION · ROADMAP	Simulated allow, deny and escalate decisions	Not certification, payment or live authorisation

Prices, where needed, come only from the current approved catalogue and are kept in separate offer sheets. An auditable event history is not a blockchain or an operational trust network.

Five questions to ask any AI supplier about status

- [ ] Is it generally available, in preview or on the roadmap?
- [ ] For which plan, region and date does that apply?
- [ ] What evidence of operation exists: customers, uptime, incidents?
- [ ] What changes with notice, and how much notice?
- [ ] What happens to our data and configuration at exit?
- [ ] Who supports it, and on which hours and terms?

ROADMAP 1 · FIRST 90 DAYS

From conclusions to a first measured decision

The six conclusions of this study translate into one 90-day sequence, counted from your own start date. Roles are proposed ownership, not evidence that people have been assigned.

1 · Connect adoption with outcomes
Track accepted outcomes, quality and full cost together.

2 · Build capability around a useful first task
One process, a baseline, time to learn.

3 · Adapt the route to the organisation
Proportionate entry for small firms; clear ownership in large ones.

4 · Use published evidence
Treat missing documentation as an open question; ask in writing.

5 · Make governance part of progress
An AI inventory with owners and the earliest deadline.

6 · Invest in learning and resilience
Staged pilots, tested fallbacks, exit options.

HORIZON	OUTPUT	OWNER	DEPENDENCY	MEASURE	DECISION
Days 1–30	One material process chosen; baseline measured; supplier, data and infrastructure dependencies mapped	Strategy lead with process owner	Access to process data	Baseline signed off	Proceed to comparison, or choose another process
Days 31–60	Evidence-led comparison of three options; one bounded pilot under a written charter; price, region and provider change tested	Process owner	Written provider facts	Accepted outcomes, quality, cost to date	Continue, adjust scope or stop early
Days 61–90	Accepted results and full cost reviewed against the threshold	Executive sponsor	Complete pilot record	Cost per accepted outcome vs baseline	Expand, redesign or stop, with a named owner

Bring to day 1

- [ ] Process named, with its accountable owner
- [ ] Budget ceiling for the pilot, incl. integration and review
- [ ] Employee representatives informed where required
- [ ] Three months of process data, or a plan to collect it
- [ ] Data-protection and security contacts
- [ ] List of AI tools already in use in that process

© 2026 ETHORITY & STEN.AI LLC · Put AI to Work: The Good AI Prism Global Market Study 2026/27 · **ethority.ai/ai-prism-study**

ROADMAP 2 · 2027

Capability progression in 2027

Planning horizons, not promises: each step is conditional on the evidence from the step before and on available resources. Regulatory dates are fixed; market conditions are not.

	H1 2027	H2 2027	TOWARD 2028
Your organisation	Scale what met its threshold; cost per accepted outcome tracked monthly	Second and third use cases; shared review and fallback routines	AI estate reviewed yearly; retire what no longer pays
Rules and dates	US state duties from 1 Jan; Vietnam grace ends 1 Mar; India 13 May	EU Annex III high-risk duties from 2 Dec	EU Annex I and Machinery: 20 Jan 2027 (machinery) and 2 Aug 2028 (AI in regulated products)
Market and infrastructure	Memory and accelerator supply tight; capacity by region varies	New model generations and lower inference prices change options	EU AI gigafactories selected in 2027; capacity no earlier than 2028
ETHORITY research	Wave 2: same provider sources re-read; changes dated	Extension of depth research beyond Customer Service and Manufacturing	Configuration-level evidence (product × plan × region)

What would change this plan: indicators to watch

INDICATOR	LATEST OBSERVATION	RESPONSE IN YOUR PLAN
Model generations and deprecation	Top-tier models priced at US$4–10 per million input tokens (Gemini 4 Argon, Claude Fable 5.1, GPT-6 Astra; Sep–Oct 2026). OpenAI shuts down its GPT-4o and 4.1 snapshots on 14 Jan 2027.	Keep integrations portable; test a successor model before a notice period ends.
Multi-model use	81% of large enterprises use three or more model families (a16z, n=100, Jan 2026).	Evaluate models per task; record which model serves which use.
Accelerators and memory	NVIDIA data-centre revenue US$89.0 bn in one quarter (+117%). ASML expects 2027–28 supply to fall short; SK hynix calls 2027 its tightest supply year.	Confirm capacity, region and price terms; watch escalator clauses.
Inference price	Prices for the same capability fall by 9× to 900× a year, depending on the task (Epoch AI).	Re-price usage contracts; the cost risk sits in volume growth.
Energy and grids	Data centres used 415 TWh in 2024, heading for about 945 TWh in 2030; about 20% of planned projects risk grid delays (IEA).	Ask where workloads run and what the fallback region is.
Physical AI	More than 600,000 industrial robots installed in 2025 (+11%), 5 million in operation; about 7,000 humanoids sold (IFR, Sep 2026).	Treat humanoids as pilots; plan for the EU Machinery Regulation from 20 Jan 2027.

Sources: provider pages and filings, a16z, Epoch AI, IEA, IFR, as recorded in the ETHORITY technology fact packs of 22 Oct 2026. Q3 2026 hyperscaler results follow at the end of October.

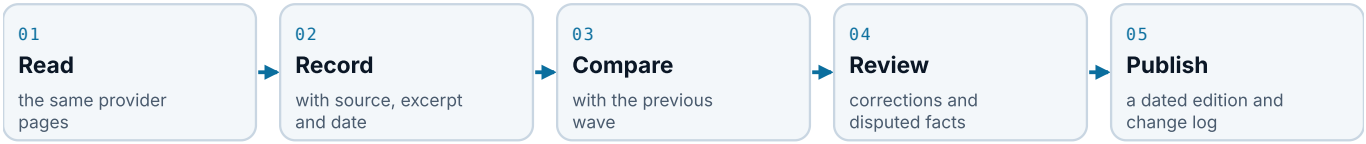
ROADMAP 3 · EVIDENCE REFRESH

How this study stays true

This is Wave 1 of a longitudinal programme. Each wave re-reads the same provider sources, dates every change and publishes corrections. A new wave changes the evidence cutoff; a new publication version does not.

STEP	WHAT CHANGES	HOW IT IS HANDLED	REVIEWER	OUTPUT
Correction window	Provider and reader submissions on this edition	Logged by entry; confirmed corrections dated	Research owner	Correction log with the next edition
Wave 2	Provider pages re-read for the same entries	Each change dated; Wave 1 kept unchanged	Methods reviewer	Change report and new snapshot
Sector and theme extensions	New maps (e.g. care, pharma, finance)	Admitted only through a dated snapshot with an overlap bridge	Research owner	Admission record; no silent totals
Technology ring	Models, chips and compute, robotics	Fact packs refreshed at each provider reporting cycle	Domain reviewer	Dated technology briefs
Method versions	Definitions, units, evidence steps	Versioned change with a bridge from old to new	Methods reviewer	Method note in the edition
Release record	Edition ID, files, checksums	One version across file, cover, footer and register	Editorial owner	Checksums listed at release

Corrections: ethority.ai/ai-prism-study. Every named provider is offered a correction window before final release. Historical editions remain available; nothing is overwritten.



What a reader can rely on

- Every figure keeps its source, date, population and definition.
- Waves are compared on the same units; nothing is silently re-based.
- Negative results and missing evidence are reported, not dropped.
- Earlier editions remain available with their own identifiers.

ROADMAP 4 · SCENARIOS AND GATES

Decisions that hold when conditions change

One test belongs in every AI decision: will it still work if prices, providers, capacity, regulation or business priorities change? The table turns each source of uncertainty into a buying response and the evidence to request.

UNCERTAINTY	BUYING RESPONSE	EVIDENCE TO REQUEST
Agenda changes	Staged commitments with scheduled reconsideration	Decision owner, trigger and review date
Provider or model change	Portability and a fallback	Export formats, switching cost, notice periods, tested substitute
Geopolitical restrictions	Jurisdiction and dependency review	Deployment region, legal supplier, subprocessors, access restrictions
Compute or energy constraints	Capacity and service review	Availability in your region, quotas, SLA and support boundaries
Financing or market correction	Continuity review	Public financial disclosures, contract protections, exit plan
Workforce resistance	Participatory change plan	Affected roles, training, oversight, plan for released capacity
Uncertain return	Measured, reversible pilot	Baseline, full cost, quality threshold, stop/go, opportunity cost

Three scenarios for 2027

SCENARIO 1

Prices fall, capability jumps

Inference becomes much cheaper and new models change what is possible.

No-regret move. Compare price certainty with flexibility; keep the integration portable.

SCENARIO 2

Supply or jurisdiction shock

Access to a provider, region or component is restricted or delayed.

No-regret move. Know your fallback and where your data and workloads run.

SCENARIO 3

Demand weakens, budgets reset

Returns disappoint and investment is cut back.

No-regret move. Stage investment with measurable pilots and a continuity plan.

Scenarios, not predictions; no probabilities are assigned. The scenario worksheet is in the evidence annex.

Questions to assess by region

REGION	WHAT TO ASSESS BEFORE COMMITTING
Europe	Regional capacity, data jurisdiction, exit options and operating cost. EU hosting alone is not complete sovereignty: ask about the legal supplier, subprocessors and support access.
United States	Provider concentration, contract terms, capacity availability by region and state-specific duties for automated decisions from 2027.
China and East Asia	Access and export controls, ecosystem compatibility and local deployment rules; confirm each against the primary source for the country concerned.
Gulf states	Power, cooling and water, delivery milestones and legal and data boundaries; an announced campus is not a commissioned one.
Emerging economies	Affordable access, connectivity, skills and practical deployment modes, more than rankings.

Assessment prompts for each region; they are not new regional findings.

WHAT THIS SHOWS

The sources of uncertainty differ by region, but the buying responses repeat: staged commitments, portability, documented jurisdiction and a measured pilot.

WHAT IT MEANS FOR YOU

Resilience is cheapest when it is written into the first contract and the first pilot charter, not added after a disruption.

WHAT TO DECIDE NOW

Pick your most critical AI use and complete the scenario worksheet for it before the next renewal or purchase.

© 2026 ETHORITY & STEN.AI LLC · Put AI to Work: The Good AI Prism Global Market Study 2026/27 · **ethority.ai/ai-prism-study**

WORKSHEET 1 · DECISION AND EVIDENCE

Buyer worksheet

One page per decision. Fill it in before the first supplier conversation; it turns this study into a decision your organisation can defend.

FIELD	YOUR ANSWER
Process and owner The process to improve and the person accountable for the result	
Baseline today Volume per month · cost per case · quality measure · cycle time	
Accepted outcome The outcome you will count (resolved case, approved invoice, reviewed code change)	
Shortlist Three options, with the reason each is on the list	
Data requirements Which data each option needs, where it lives, who may share it	
Full cost Licence or usage · integration · human review · rework · change	
Decision rule The result that justifies a pilot · the result that stops it	

Provider evidence checklist

Ask each shortlisted provider for these facts in writing, with the date and the document they come from. An unanswered item is an open risk, not a pass.

- [ ] Product and plan in scope
- [ ] Use of customer data for training
- [ ] Integrations with your systems of record
- [ ] Certifications and their scope
- [ ] Human oversight and escalation
- [ ] Hosting region and subprocessors
- [ ] Retention and deletion
- [ ] Billing unit and what drives it
- [ ] Exit: data export and notice periods
- [ ] Incident and change notification

WORKSHEET 2 · PILOT AND OUTCOME

Pilot charter

Agree it before the pilot starts. A pilot without a threshold and a stop condition cannot fail, and so cannot prove anything.

ELEMENT	AGREED
Scope and duration Process, users, volume and the weeks the pilot runs	
Success threshold The level of accepted outcomes that justifies scaling	
Quality controls Who reviews which outputs, and how errors are recorded	
Budget ceiling Total spend the pilot may not exceed, including integration and review	
Stop conditions The results or incidents that end the pilot early	
Decision date and owner When the go/stop decision is taken, and by whom	

Cost per accepted outcome

The measure that makes AI options comparable across billing units and suppliers. Use one currency, one period and the same cost boundary for the AI option and for today's process. Count every attempted case; if no outcome is accepted, the measure is undefined, not zero. Savings are a separate calculation.

Cost per accepted outcome

$$\frac{\text{usage or licence + amortised integration + monitoring, change and training + human review + handling of failed or escalated cases}}{\text{attempted cases} \times \text{acceptance rate}}$$

Worked example (illustrative)

A support team receives 10,000 cases a month. Per month, in euros: usage €6,000; amortised integration €1,500; monitoring, change and training €500; human review €2,000; handling of the 6,000 failed or escalated cases €1,000 (rework counted here only, not twice). The tool resolves 4,000 cases (40%) to the customer's acceptance.
€11,000 ÷ (10,000 × 0.40) = €2.75 per accepted outcome, to compare with today's cost per resolved case on the same monthly boundary.

Evidence ledger

Keep one line for every number you rely on: the claim, the exact source, publication and access dates, the population, the definition and the limitations. This study follows the same rule; corrections and exclusions are listed by edition.

ABOUT THIS STUDY

Method and limitations

The study combines public market and buyer research published in 2025 and 2026 (official statistics, industry associations, academic studies and consulting surveys, each cited where it is used) with the Good AI Prism research edition of 20 October 2026. Survey figures are not pooled: samples, questions and definitions differ. The Prism contains 998 market entries, including 961 AI offerings, across 24 business areas, 8 data types and 6 technology layers; each entry has a dated provider source. The depth-research set contains 194 records by unique research ID across Customer Service and Manufacturing, including 12 earlier Customer Service profiles: pricing unit, published price, hosting, use of customer data for training, integrations and limitations. Record counts do not establish independent product uniqueness or verified performance.

The global comparison uses official statistics wherever they exist (Eurostat and national statistics offices, IMDA, ABS, IBGE, Banxico), labels vendor and sponsored surveys as such, and does not add up regional spending forecasts that come from different releases and definitions. More than 30 economies are covered; for several, including the UAE, Saudi Arabia, India and China, no official firm-level adoption statistic was available.

Research was carried out with AI tools under human direction. A random sample of recorded facts was re-checked against the cited pages in a separate review. ETHORITY reviewed the provider pages behind every finding, and every named provider can submit corrections. The authors retain editorial responsibility.

A longitudinal research programme. This study is Wave 1 of ETHORITY's longitudinal research programme. The programme observes the same provider sources repeatedly: in each wave ETHORITY reads the providers' own pages, records what they document with source and date, and compares the result with the previous wave. Wave 1 documents the state of the provider information ETHORITY reviewed on 19–20 October 2026. Changes are recorded in subsequent waves. The market and buyer figures are a separate synthesis of external studies from 2025 and 2026; each is cited with its own source, sample and method and is not pooled.

Reading the figures

- **Interpretation is marked.** Under each exhibit, "What this shows" restates the evidence; "What it means" and "What to decide now" are the authors' interpretation and proposals.
- **Survey answers are not a funnel.** Shares of organisations that use, scale or profit from AI come from separate questions; they are not shown as stages through which the same firms pass.
- **Not legal advice.** References to the AI Act and the Digital Omnibus describe the public record as of October 2026.
- **Surveys measure different things.** "Any use", "regular use" and "production use" are not comparable. Each figure is shown with its source, sample and date. Where a figure comes from a secondary or vendor-sponsored source, the row says so.
- **Germany by firm size.** ifo shows medium firms (47.2%) below small ones (51.2%), but does not define its size classes. The Eurostat 2025 split by size (about 23% vs 57%) comes from a secondary analysis.
- **Channels are inferred.** No German survey we found measures system houses or IHKs as buying channels; the role of these channels is our judgement.

Counts and denominators

- **998 market entries** in the research edition of 20 October 2026: 961 AI offerings (641 products, 306 products named after their provider, 14 products identified by provider), 33 provider records and 4 entries awaiting classification. They hold 1,037 placements on the wheel.
- **194 depth-research records**, counted by unique research ID: 182 in this edition (95 in Customer Service, 87 in Manufacturing), plus 12 Customer Service profiles researched earlier. The Customer Service cohort is therefore 107 records. Records are not a verified count of distinct commercial products; possible overlaps are documented separately.
- **10 of 107**: recomputed by offering ID from the frozen fact matrix: 38 document an EU hosting option, 23 document limits on training (20 no training, 3 opt-in only), 10 document both (8 no training, 2 opt-in only).

Taxonomy

The Prism's 24 business areas are 14 functions and 10 industries. The 11 sectors in chapter 9 are a reading of those industries plus legal services, which the Prism maps as a function (Legal & Compliance). Telecoms, construction and real estate, professional services and hospitality have no industry petal yet; sector briefs for them will be assembled from the functions they use.

© 2026 ETHORITY & STEN.AI LLC · Put AI to Work: The Good AI Prism Global Market Study 2026/27 · **ethority.ai/ai-prism-study**

SOURCES, CREDITS AND AUTHORS

Who made this study, and where the evidence lives

Created by Sten Franke, with a people and governance lens by Severin L. Wilson. Every figure links to its source in the chapter where it is used; the complete register is published with the study.

Sten Franke
Founder, Principal Architect & Lead Creator. Creator and architect of ETHORITY's AI systems, the Good AI Prism and the AI Business Operating System; builds companies since 1997 (mediatime, Gridpatrol, ETHORITY) and created the Social Media Prisma.
linkedin.com/in/stenfranke · sten@ethority.ai

Severin L. Wilson
Chief Business & Transformation Officer · Human Capital, AI Transformation & Governance. Connects business ambition with the people, capabilities and governance needed to deliver change; experience involving Detecon, SAP SuccessFactors and E.ON.
linkedin.com/in/severinwilson · severin@ethority.ai

Key sources

Gartner, AI spending forecast, Sep 2026
McKinsey, The State of AI, Aug 2026
Microsoft, AI Diffusion Report 2026
Eurostat, ICT usage in enterprises 2025
ONS, INSEE, Istat, INE, CBS, ABS, IMDA, IBGE, Banxico (national statistics)
Bitkom, KI-Studie, Sep 2026; Bitkom Cloud Report 2026
IDC regional AI spending guides, 2024–2026
Stanford AI Index 2026
EUR-Lex, Regulation (EU) 2026/1744
National and law-firm sources for 17 jurisdictions (evidence annex)
NBER working paper 31161 (customer support)
METR and Google RCT (software development)
IEA, Energy and AI; Key Questions on Energy and AI, 2026
IMF, Global Financial Stability Report, Oct 2026
WEF Global Risks Report; Allianz Risk Barometer; PwC CEO Survey, 2026
Good AI Prism research edition, 20 Oct 2026 (998 entries; 194 records)

Credits and rights

- Good AI Prism artwork and all exhibits: © 2026 ETHORITY & STEN.AI LLC. Exhibits may be quoted with source attribution.
- Data remain the property of the cited publishers; figures are reproduced with attribution.
- Research carried out with AI tools under human direction; the authors retain editorial responsibility.

Edition

Research edition 2.0 (40 pages), 22 October 2026, review draft; evidence cutoff unchanged from Wave 1 (provider pages read 19–20 October 2026). Builds on the first edition v1.9. Corrections: ethority.ai/ai-prism-study.

Evidence annex and complete source register: ethority.ai/ai-prism-study

About ETHORITY. ETHORITY connects market intelligence, research and practical delivery to help organisations put AI to work. Its roots lie in the digital intelligence business Sten Franke founded in 2005; ETHORITY & STEN.AI LLC was established in 2023. It works as an AI1st organisation: people and AI agents on one AI Business Operating System, with the research, engineering and operations to build and run AI after the pilot.

ETHORITY
AI BUSINESS OPERATING SYSTEM



FROM BUSINESS AMBITION TO AI THAT WORKS

Put AI to work, with evidence.

ETHORITY helps organisations choose, govern and prove AI: from independent market evidence to decisions that can be defended and outcomes that can be measured.

CONTACT



Sten Franke
Founder, Principal Architect & Lead Creator
sten@ethority.ai
linkedin.com/in/stenfranke



Severin L. Wilson
Chief Business & Transformation Officer
severin@ethority.ai
linkedin.com/in/severinwilson

STUDY REPOSITORY AND UPDATES

ethority.ai/ai-prism-study

ethority.ai

THE ETHORITY STUDY SUITE

01	**Put AI to Work: The Good AI Prism Global Market Study 2026/27**	Global edition · EN and DE · YOU ARE READING
02	**Put AI to Work: The Good AI Prism Market Study 2026/27**	Germany–US focus · EN and DE
03	**Good AI Prism sector, function and technology briefs 2026/27**	HR · Health · Financial Services · Public Sector · Education · Energy · Mobility · Finance & Accounting · AI Models · Chips & Compute · Robotics · EN and DE
04	**AI for Care in Ageing Societies**	Study 2026/27 · EN, DE, US and Global editions
05	**AI in Pharma, Biotech and Longevity**	Study 2026/27 · Global EN and DE
06	**The Social Media Prisma 2009–2025**	Historic paper · EN and DE

© 2026 ETHORITY & STEN.AI LLC · San Francisco HQ · European operations: Hamburg, Germany · ethority.ai
Put AI to Work: The Good AI Prism Global Market Study 2026/27 · Research edition 2.0 (40 pages), review draft, 22 October 2026 · All rights reserved. Exhibits may be quoted with source attribution.