ETHORITY®
AI BUSINESS OPERATING SYSTEM

FIRST EDITION · WAVE 1 · OCTOBER 2026

RESEARCH STUDY · GERMANY & UNITED STATES FOCUS 2026/27

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

Germany & United States Focus

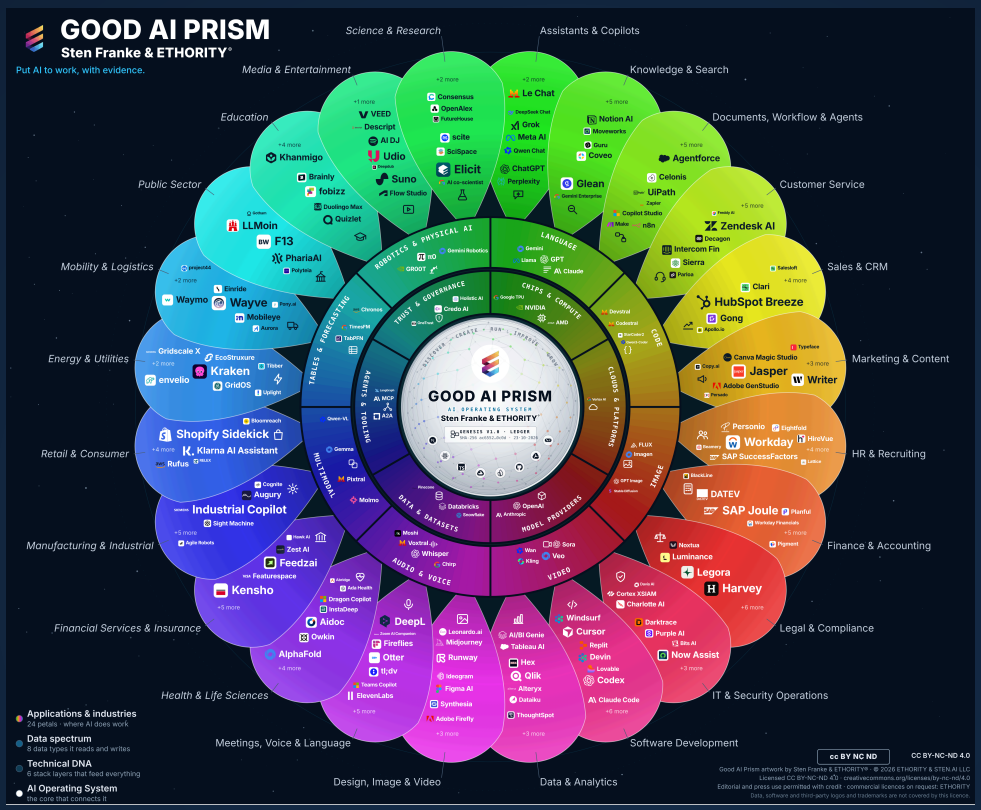
What the German Mittelstand, US businesses and enterprises need from AI now: sector by sector, function by function, what the market's own documentation reveals, and the decisions that will shape 2027.

961
AI OFFERINGS MAPPED

24
BUSINESS AREAS

11
SECTORS ANALYSED

10 of 107
DOCUMENT BOTH KEY DATA FACTS · WAVE 1



Artwork: v3.11, reference selection of 17 October 2026. Statistics: separate research inventory, 20 October 2026.

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

Source-based market research
Germany · United States · sectors

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?* Chapter 8 shows 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 practical appendix turns them into tools you can use on one decision.

CONTENTS

How this study is built

From the market to your next decision: the evidence first, then what it means, then tools to act on it.

OVERVIEW

Founder's perspective 2
In 90 seconds 4
The Good AI Prism 5
Ten findings 6

THE EVIDENCE

01 The market in 2026 7
02 Regulation 8
03 Germany 9
04 United States 11
05 Sectors 12
06 Business functions 13
07 What the market documents 15

WHAT IT MEANS

08 AI decisions in an uncertain world 16
Conclusions 18
09 The 2027 agenda 19
A second lens: work and people 20
10 What leaders should do next 21

TOOLS AND BACKGROUND

A Buyer worksheet 22
B Pilot charter 23
C Workforce annex 24
D Scenario worksheet 25
Method and evidence 26
About the authors 29
About ETHORITY 30
Sources 32

IN 90 SECONDS

AI adoption is advancing unevenly. Organisations need practical ways to select useful applications, build capability and measure results. In 2026, 57% of German firms with 20 or more employees use AI, yet across the US and the EU only about one firm in five uses it to produce goods or services, and 37% of organisations in a global survey attribute some EBIT effect to AI (McKinsey, Aug 2026). The constraint is no longer access to tools. It is know-how, comparable facts, a cost case and a decision that can be defended. Our review of 107 customer-service AI offerings shows how hard that 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

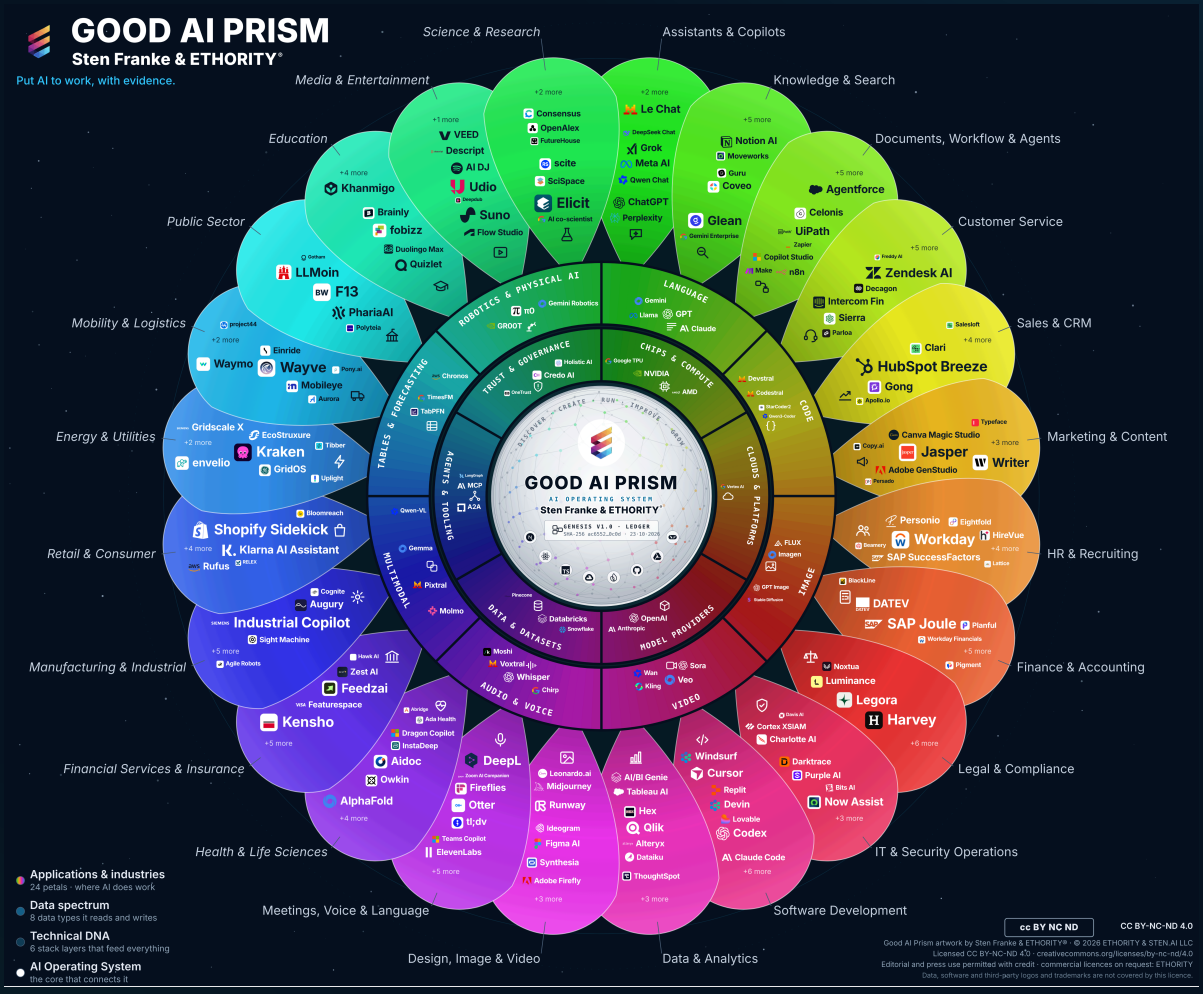
QUESTION	FINDING	WHAT LEADERS SHOULD DO
How far has adoption come?	It differs widely: 57% of German firms with 20+ employees and close to 90% of larger organisations report use, while about one firm in five uses AI to produce goods or services.	Move from tool trials to two or three measured use cases.
What holds the Mittelstand back?	Know-how (85% of non-users), legal uncertainty (66%) and costs that are hard to calculate (63%).	Start with a readiness assessment and a cost case per process.
What do US businesses need?	Help choosing (48% of small-business users struggle), proof of profit (37% attribute some EBIT effect to AI) and control of usage costs (36% of large firms have it).	Choose by job and total cost; put usage-based spend under control.
Is regulation a blocker?	It is a timetable by covered use: EU Annex III high-risk duties from 2 December 2027, Annex I product-related systems from 2 August 2028; covered automated decisions in Colorado and California from 1 January 2027.	Build the AI inventory and governance now, ahead of the dates.
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 different 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.

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.

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 Germany and the United States.

01 **Use is broad; scale is rare.**
"Any use" among larger firms is near 90% in the US and EU, but production use across all firms is about one in five.

02 **Germany crossed the majority line in 2026.**
57% of firms with 20+ employees use AI, up from 36% a year earlier (Bitkom). Most still use it for text, not processes.

03 **In Germany the gap is know-how, not access.**
85% of non-users lack technical know-how; 66% rate their staff's AI skills as low.

04 **German firms can't yet calculate the cost.**
Cloud, data preparation and integration drive cost far more than licences (51%, 50%, 41% vs 21%), and 64% of small firms can't judge the benefit.

05 **US small firms need help choosing; enterprises need cost control.**
Among US small-business AI users, 48% struggle to choose the right tool and 73% want help to implement and evaluate AI (Goldman Sachs 10KSB). Only 37% of organisations attribute some EBIT effect to AI (McKinsey, Aug 2026), and only 36% of large US firms have direct controls on token and usage costs (KPMG).

06 **Sovereignty is now a buying criterion.**
75% of German firms call Germany too dependent on US cloud providers.

07 **Governance trails strategy.**
Of German large firms, 98% have an AI strategy and 37% have clear governance. The AI Act high-risk dates are now fixed at 2 Dec 2027 and 2 Aug 2028.

08 **Across the sectors we reviewed: which AI, how safely, does it pay?**
Energy terminates 34% of AI initiatives; only 7% of logistics firms trust AI output highly; only 18% of legal firms track ROI.

09 **Providers rarely publish the facts buyers need.**
In 10 of 107 customer-service AI offerings reviewed, 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).

10 **Breadth is easy to find; verified depth is scarce.**
Directories list 30,000 to 50,000 AI tools. Sourced, buyer-relevant facts with comparable alternatives are what decisions need, and they are rare.

1 · THE MARKET IN 2026

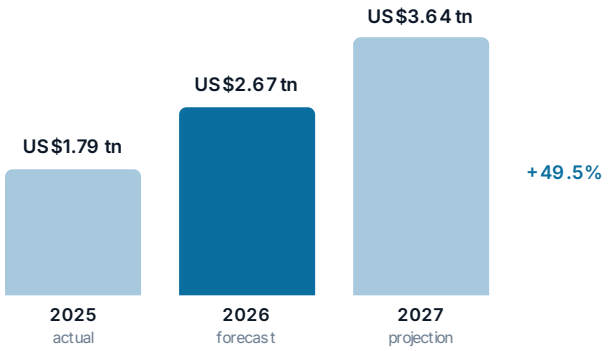
Big spend, broad use, rare scale

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 · REGULATION

The clock is set, on both sides of the Atlantic

The EU AI Act sets the clock for Europe. The Digital Omnibus on AI, Regulation (EU) 2026/1744, has applied since 27 July 2026. It moves high-risk obligations under Annex III to 2 December 2027 and AI in regulated products (Annex I) to 2 August 2028. Transparency duties under Article 50 apply from 2 August 2026; generative AI already on the market has until 2 December 2026 for machine-readable marking. The AI literacy duty in Article 4 now requires measures to support staff literacy. In the US there is no federal AI act, but state rules carry dates.

EXHIBIT 3

The EU AI Act timetable after the Digital Omnibus

Application dates under Regulation (EU) 2026/1744, in force since 27 July 2026



2 Aug 2026
Article 50 transparency duties apply

2 Dec 2026
Machine-readable marking for generative AI already on the market (Art. 50(2))

2 Dec 2027
High-risk obligations for uses listed in Annex III, e.g. in employment, credit, insurance, education and public services

2 Aug 2028
High-risk AI in regulated products, Annex I; some public-sector systems 2 Aug 2030

Source: EUR-Lex, Regulation (EU) 2026/1744; dates cross-checked against the Council text PE-CONS 30/26. Whether a system is high-risk depends on its intended use. This study is not legal advice.

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

United States: dated state rules, no federal AI act

JURISDICTION	WHAT APPLIES
Colorado	SB 26-189, signed 14 May 2026, repeals and re-enacts the 2024 Colorado AI Act as a narrower law on automated decisions. From 1 January 2027 deployers must give notice before use, explain an adverse consequential decision within 30 days, offer correction and human review where commercially reasonable, and keep records for three years (Colorado session law; Eckert Seamans).
California	CCPA regulations effective 1 January 2026 (CPPA); businesses using automated decision-making technology for significant decisions must comply from 1 January 2027, or before first use if adopted later (Hunton).

Small US firms see regulation mainly as cost: 65% worry about a patchwork of state AI laws (U.S. Chamber).

Reading the dates

Application dates differ from enactment dates. Map each system to its jurisdiction, covered use and applicable duty. Colorado and California requirements for covered automated decisions apply from 1 January 2027. EU Annex III high-risk uses follow from 2 December 2027 and Annex I product-related systems from 2 August 2028. The EU Omnibus changes high-risk dates; existing transparency and AI literacy duties remain relevant.

WHAT THIS SHOWS

Rules differ by model and jurisdiction, but every major regime now has dated obligations, and labelling of AI-generated content is common to most.

WHAT IT MEANS FOR YOU

Regulation is a schedule to plan against, not a reason to wait. The organisations that know what AI they run will meet the dates at the lowest cost.

WHAT TO DECIDE NOW

Build one inventory of the AI systems in use, record where each is deployed, and map it to the earliest deadline that applies.

3 · GERMANY

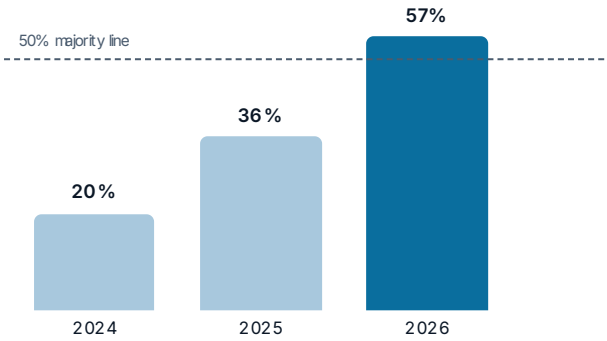
What the Mittelstand and large enterprises need

German firms crossed the majority line in 2026. Most still use AI for text work, though, not for processes. The constraint is no longer access to tools. It is know-how, clarity on the law and a cost case firms can calculate.

EXHIBIT 4

AI crossed the majority line

German firms with 20+ employees using AI



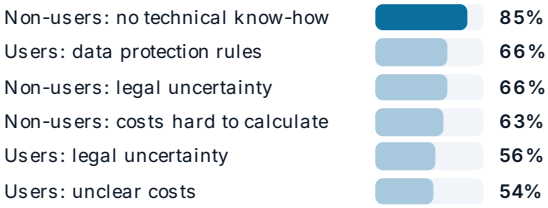
Source: Bitkom, Sep 2026 (n = 603). Base differs from Eurostat (10+ staff).

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

EXHIBIT 5

Know-how, law and cost hold firms back

Share of firms naming each barrier



Source: Bitkom, Sep 2026

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

INDICATOR	VALUE	SOURCE
Firms (20+ employees) using AI	57% (2025: 36%, 2024: 20%)	Bitkom, Sep 2026
Use by firm size	Large 67.2%, small 51.2%, medium 47.2% (54.5% overall)	ifo, Jun 2026
Mittelstand using AI (2022–24)	20%, about 780,000 firms; text generation 14%, automation and decision support only 5%	KfW Fokus 554
German AI market	€28.7 bn in 2026 (+48%), €40.3 bn forecast for 2027	Bitkom/IDC
Firms that see their AI potential as fully used	0%; 59% say "not at all"	DIHK 2026
Firms that call Germany too dependent on US cloud providers	75%; 98% say a provider's country of origin matters	Bitkom Cloud Report 2026
Large firms: AI strategy and governance	98% have a strategy; 37% have clear governance	KPMG, Jun 2026
AI "leaders"	6% in Germany, 15% globally	BCG AI Radar 2026

REGIONAL CONCLUSION · GERMANY

OBSERVED CONSTRAINT

85% of non-users lack know-how; 63% find costs hard to calculate; in the Mittelstand, text generation (14%) far outweighs automation and decision support (5%).

OUR INTERPRETATION

German adoption is broad but shallow. Large firms have strategies (98%) but rarely governance (37%); sovereignty (75%) is becoming a buying criterion.

NEXT DECISION

Mittelstand: one process, one owner, the full cost written down. Large enterprises: an inventory of AI in use with a named owner for each system.

2027 IMPLICATION

High-risk dates and sovereignty expectations will shape supplier choices in 2027; firms without an inventory will decide under time pressure.

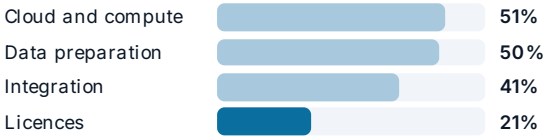
What the Mittelstand needs

BUSINESS PROBLEM	EVIDENCE	WHAT HELPS	OUTCOME TO MEASURE
No know-how	85% of non-users lack technical know-how; 66% rate staff AI skills as low (Bitkom)	Readiness assessment; staff AI literacy training	First process live and reviewed at an agreed horizon
A cost case they can't calculate	63% of non-users find costs hard to calculate; 64% of small firms can't judge the benefit (IW)	Compare pricing units; business case with a stop/go gate	Cost per process, payback in months
Legal and data protection	Data protection 66%, legal uncertainty 56% among users (Bitkom)	Check EU hosting and training terms before purchase; AI Act readiness review	Documented tool decisions; staff AI literacy measures
Hidden costs	Cloud/compute 51%, data preparation 50%, integration 41%; licences only 21% (Bitkom)	Assess integration effort and prerequisites; total cost of ownership	Total cost of ownership, not only the licence

EXHIBIT 6

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

Channels. The Mittelstand buys AI through people it already trusts: system houses, chambers of commerce (IHKs) and the Mittelstand-Digital centres. No survey we found puts a percentage on these channels, so this is our inference. The practical route into the Mittelstand runs through these intermediaries.

What large enterprises need

BUSINESS PROBLEM	EVIDENCE	WHAT HELPS	OUTCOME TO MEASURE
Strategy without steering	Only 39% say top management actively steers AI (KPMG)	Governance with named owners and KPIs per use case	Owner and KPIs for every use case
Unknown AI inventory and shadow AI	Supplier sprawl; high-risk deadlines 2 Dec 2027 and 2 Aug 2028 (Digital Omnibus)	Complete AI inventory with risk classification, kept separate from supplier marketing	Full register; risk classes assigned
Sovereignty and lock-in	59% cite lock-in as a reason not to switch; 80% want German or EU hyperscalers (Bitkom Cloud)	Sovereignty criteria: hosting, provider origin, exit options	Share of critical AI with an EU option
Scale beyond pilots	Germany has 6% AI leaders, 15% globally (BCG)	Business case per use case with ROI tracking	Value realised per use case

4 · UNITED STATES

What US small businesses and enterprises need

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 2)	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

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.

OUR INTERPRETATION

The US market is a market of choice and return. Usage-based pricing makes cost control the next management task.

NEXT DECISION

Compare tools for one job on total cost at your volume; make usage costs visible per use case before adding agents.

2027 IMPLICATION

Colorado and California rules on automated decisions apply from 1 January 2027 for covered uses; map automated decisions now.

5 · 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.

6 · BUSINESS FUNCTIONS

High use, thin proof

Most functions we reviewed 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 anywhere from 19% slower to 21% faster in software work. The common need is the same: choose well, measure honestly, govern the risk.

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

6 · BUSINESS FUNCTIONS, CONTINUED

FUNCTION	WHERE IT STANDS	BUSINESS PROBLEM	WHAT HELPS	OUTCOME TO MEASURE
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.

7 · WHAT THE MARKET DOCUMENTS

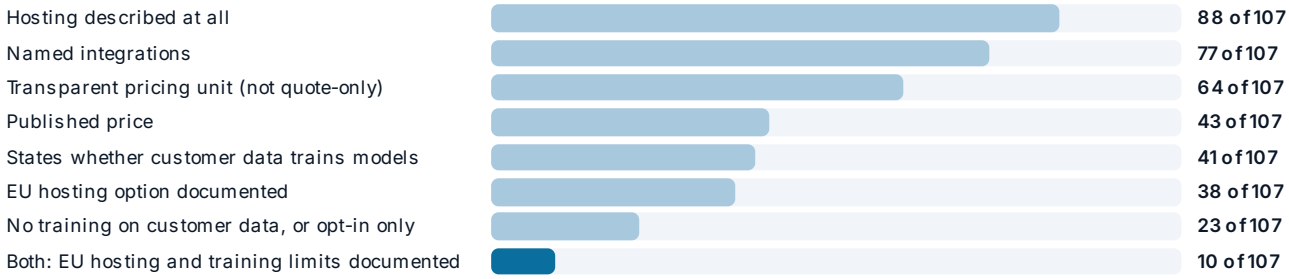
The facts a careful buyer needs are mostly not public

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 7

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.

What leaders are trying to avoid

Companies are choosing AI while the conditions for using it change: markets, rules, financing, infrastructure and the organisation's own priorities. A credible AI strategy must remain useful when those conditions change.

WEF · GLOBAL RISKS REPORT 2026
Over 1,300 global leaders and experts; risk perception

#1	Geoeconomic confrontation: the risk most likely to trigger a global crisis in 2026
18%	of respondents chose it as that most likely trigger
#1	Also first for severity over the next two years

Source: Global Risks Perception Survey, 14 Jan 2026

ALLIANZ · RISK BAROMETER 2026
Companies of all sizes; share of responses, several answers possible

42%	Cyber incidents: top risk for the fifth year, first in every region
32%	Artificial intelligence: up from 10th in 2025 to 2nd, the largest jump
29%	Business interruption, outside the top two for the first time

Source: Allianz Commercial, 14 Jan 2026

PWC · 29TH GLOBAL CEO SURVEY
4,454 CEOs in 95 countries and territories; fieldwork Sep–Nov 2025

42%	name keeping pace with technological change, including AI, as their top concern
12%	report both lower costs and higher revenue from AI
56%	report no significant financial benefit from AI so far

Source: PwC, 19 Jan 2026

Three different populations and questions, shown separately; the percentages are not comparable with one another. Survey results describe perceptions, not incident rates or outcomes. Fieldwork for all three predates this study by up to a year.

Our reading: the practical fear is committing to a system that becomes expensive, unavailable, unsuitable or strategically irrelevant before its value has been realised. The opposing fear is waiting while competitors redesign the work. Leaders need choices that create useful capability now and keep room to change direction.

Six decisions, two risks each

DECISION	RISK OF MOVING TOO FAST	RISK OF MOVING TOO SLOWLY	A REVERSIBLE WAY THROUGH
Invest now or stage it	Paying for capability that is soon cheaper or obsolete	Falling behind competitors who redesign the work	Staged commitments with a review date
Buy, build or combine	Lock-in to one supplier's roadmap	Owning maintenance the business cannot staff	Hybrid with clear exit and export terms
Keep control of data	Data used or stored beyond what is agreed	Blocking the uses that create value	Documented hosting, training and access terms
Rely on which providers	Concentration: price, capacity or access changes	Integration cost of too many suppliers	A tested fallback for each critical use
Redesign whose work	Resistance, skill gaps and unowned errors	Gains that stay as unused saved time	Participatory change plan with owners
Evidence before scale	Scaling an unproven result	Endless pilots that never decide	Stop/go criteria agreed in advance

Workforce lens (Severin L. Wilson). When agendas change quickly, the people side breaks first: half-adopted tools, skills that lag the system, employees who were not involved, and decisions no one owns. Plan adoption, training, employee participation and accountable oversight as part of every staged commitment, not after it.

8 · AI DECISIONS IN AN UNCERTAIN WORLD

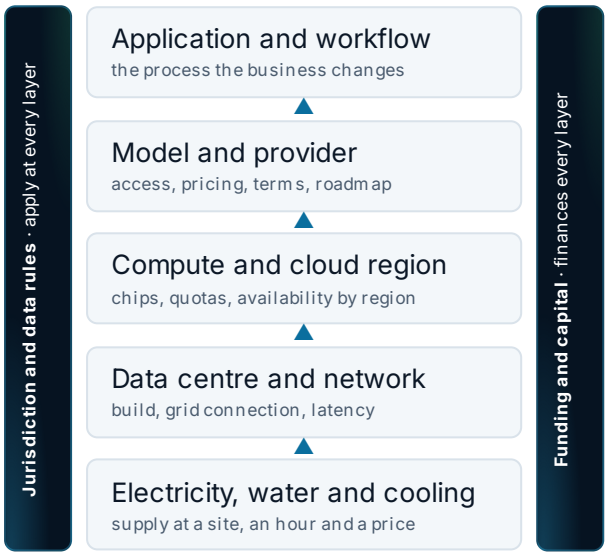
AI runs on physical and financial foundations

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 8

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 9

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.

CONCLUSIONS

What the evidence adds up to

AI creates an opportunity to improve work. These six conclusions translate the evidence into decisions organisations can take now.

1 Connect adoption with outcomes.

Evidence. Global surveys show widespread use alongside uneven financial results: about nine in ten organisations report regular use, 37% an EBIT effect. Official statistics count about one firm in five in the EU using AI.

What to do. Track adoption, accepted outcomes, quality and full cost together; do not read separate surveys as one conversion funnel.

2 Build capability around a useful first task.

Evidence. Know-how is the barrier named most often: 71% of EU firms that considered AI, 85% of German non-users, 81% of Korean manufacturers without AI specialists.

What to do. Choose one worthwhile process, measure a baseline and give people the time, skills and support to learn.

3 Adapt the route to the organisation.

Evidence. In every national study we compared, large firms report AI use 1.6 to 4.3 times as often as small ones; the gap differs by country.

What to do. Offer smaller firms proportionate entry points, reusable sector patterns and practical support; give large enterprises clear ownership across systems and functions.

4 Use published evidence to improve the shortlist.

Evidence. The research matrix records 10 of 107 Customer Service offerings with both an EU hosting option and limits on training with customer data, including two opt-in-only cases. Wave 1 documents the state of the provider information ETHORITY reviewed on 19–20 October 2026. Changes are recorded in subsequent waves.

What to do. Treat missing documentation as an open question: ask suppliers to confirm plan, region and terms in writing.

5 Make governance part of progress today.

Evidence. Duties already apply in many jurisdictions, and further rules phase in through 2027 and beyond: covered automated decisions in Colorado and California from 1 January 2027, EU duties for listed high-risk uses from 2 December 2027.

What to do. Keep an AI inventory with named owners and proportionate privacy, security, workplace and sector checks; confirm the operative law for the intended use before relying on a date.

6 Invest in learning and resilience.

Evidence. More than US$700 billion of hyperscaler investment meets grid and permitting limits; 56% of CEOs report no significant financial benefit from AI so far.

What to do. Run staged pilots with measurable hypotheses; compare price certainty with flexibility, keep tested fallbacks and exit options, and scale when the evidence supports the objective, including service quality or public benefit.

9 · AGENDA

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 10

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 11

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 the practical appendix). 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.

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. The workforce annex to the pilot charter (practical appendix C) turns these questions into a form.

10 · IMPLICATIONS

What leaders should do next

The evidence points to the same sequence for every type of organisation: decide on facts, start small and measured, govern before the deadlines, and keep what meets its objectives.

Founders and solopreneurs

- Pick one recurring task, not a platform.
- Compare three tools on cost per unit at your volume and on data terms.
- Keep a tool if it meets its objective at your volume; review it at a short, agreed horizon.

Mittelstand

- Start with a readiness assessment and staff AI literacy.
- Put one process live, track cost per accepted outcome and review it at an agreed horizon.
- Classify every AI use under the AI Act before December 2027.

Large enterprises

- Build a complete AI inventory, including shadow AI.
- Name an owner and KPIs for each use case; add stop/go gates.
- Control usage-based costs and set sovereignty criteria.

Public sector, science and education

- Set sovereignty and openness criteria before procurement.
- Ask for procurement-ready evidence, not demos.
- Publish an AI policy; train staff and students.

AI providers

- Publish pricing units, hosting regions and data-use terms.
- Documentation is becoming a buying criterion: 97 of 107 offerings leave at least one key fact undocumented.

Boards and investors

- Ask for outcomes, not usage: 37% report any profit effect.
- Treat AI governance as a 2027 deadline, not a future project.

PRACTICAL APPENDIX · A

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

PRACTICAL APPENDIX · B

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.

PRACTICAL APPENDIX · C

Workforce annex to the pilot charter

Complete it with the pilot charter. It makes the people side of the decision as concrete as the technical and financial side.

ELEMENT	AGREED
Affected roles Which roles and how many people the pilot touches	
Baseline workload Volume, time and quality of the work today, measured before the pilot	
Skills needed Operating, reviewing and improving the system; training plan and owner	
Employee consultation Works council or staff representatives where applicable; when and how they are involved	
Oversight owner The named person accountable for outputs that affect people	
Escalation and appeal How employees and affected people can challenge an output, and the response time	
Capacity plan Where released time goes: service, quality, growth or redeployment; who decides	
Measured quality Error rate, complaints and review findings, tracked with the outcome measure	

Saved time is not realised value. Count the value only when released capacity has demonstrably gone into service, quality or growth, using the same boundary as the cost-per-accepted-outcome calculation.

PRACTICAL APPENDIX · D

Scenario worksheet

Use it for one AI decision. Fill in each column before the commitment and revisit it when an indicator moves.

ELEMENT	1 · PRICES FALL, CAPABILITY JUMPS	2 · SUPPLY OR JURISDICTION SHOCK	3 · DEMAND WEAKENS, BUDGETS RESET
Our assumptions in this scenario			
Early indicators we will watch			
What changes for our AI use case			
No-regret action now			
Trigger to revisit the decision			
Owner and review date			

These are scenarios, not predictions. Do not assign probabilities you cannot support; record the assumptions instead, so that the decision can be reviewed against what actually happens.

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.

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 5 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.

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.

APPENDIX

The ten offerings behind Exhibit 7

- **Cohort:** all 107 Customer Service offerings with decision research, researched in two waves in 2026. Inclusion was by buyer job (chat, voice, agent assist, knowledge and self-service, omnichannel), not by size or popularity.
- **Criterion 1, EU hosting:** a first-party page states that customer data can be hosted in the EU. Values such as "region unclear", "UK only", "EU for models only" or a self-host option do not count.
- **Criterion 2, training:** a first-party page states that customer data is not used to train the provider's models, or is used only with opt-in consent. "Own use not stated" does not count.
- **Reading the result:** 97 offerings lack public documentation for at least one criterion. That is a finding about what providers publish, not about what their products can do.

OFFERING	TRAINING ON CUSTOMER DATA	FIRST-PARTY SOURCES
BOTfriends X	Not used (stated)	GDPR page
Dynamics 365 Customer Service	Only with opt-in	Data and privacy FAQ, geographic availability
eesel AI	Not used (stated)	Security and privacy
Kauz	Not used (stated)	Customer service page
Observe.AI	Not used (stated)	Trust centre, LLM information
PolyAI	Only with opt-in	Compliance, security
Replicant	Not used (stated)	Safety and security
sipgate AI Agents (Flow)	Not used (stated)	Data protection FAQ
Superchat AI Agent	Not used (stated)	Security
Typewise	Not used (stated)	Security

Listed alphabetically. Every named provider is offered a correction window before final release.

ABOUT THE AUTHORS

Created by Sten Franke, with Severin L. Wilson



Sten Franke

Founder, Principal Architect & Lead Creator

THE CREATOR BEHIND THE STUDY

Sten Franke is a digital pioneer, entrepreneur and AI systems architect who has built companies since 1997. Through mediatime, Gridpatrol and ETHORITY, his work has served hundreds of national and international brands and public institutions. His contributions include the Social Media Prisma and industry work involving BITKOM and BVDW.

He creates ETHORITY's vision, products and research methods, including the Good AI Prism and AI Business Operating System. His work also extends into biomedical innovation: from advising the Dr. Rurainski Foundation and developing its digital presence to EpiTrace.AI and an agenda focused on cancer research, longevity and healing. Across these fields, he connects market intelligence, engineering and entrepreneurship to build useful AI systems with clear economics, evidence and human accountability.

Sten Franke on LinkedIn



Severin L. Wilson

Chief Business & Transformation Officer · Human Capital, AI Transformation & Governance

PEOPLE TURN CAPABILITY INTO VALUE

Severin L. Wilson connects business ambition with the people, capabilities and governance needed to deliver change. Drawing on experience involving Detecon, SAP SuccessFactors and E.ON, he brings expertise in human capital management, digital transformation and international organisational change.

At ETHORITY, his focus is developing client relationships and translating AI opportunities into organisational progress: aligning leadership, developing workforce capabilities, involving employees and establishing accountable human oversight. His contribution to the study examines how organisations turn technical capability into improvements they can sustain.

Severin L. Wilson on LinkedIn

OUR PURPOSE

Help organisations make progress with AI.

Start with a useful business or public-service objective. Make room for learning, involve the people affected, and build confidence through evidence. This study offers a route from an opportunity to a practical next step, with responsibilities and outcomes that can be understood and reviewed.

ABOUT ETHORITY

From understanding markets to building AI capability

ETHORITY connects market intelligence, research and practical delivery to help organisations put AI to work. ETHORITY is built as an AI1st organisation: humans and AI agents work on one AI Business Operating System, with the research, engineering and operations to build and run AI after the pilot.

Its roots lie in the digital intelligence business Sten Franke founded in 2005. As search, social platforms and online communities changed the relationship between companies and their markets, ethority combined monitoring, research and advisory work to help organisations interpret that change and respond.

ETHORITY & STEN.AI LLC was established in 2023. It builds on that heritage while addressing a larger question: how can an organisation turn rapidly changing AI capabilities into useful, accountable work?

The Good AI Prism is the entry into the market: research profiles, documented relationships and comparisons help users investigate the options. Advisory, governance and transformation work connect those options with an organisation's objectives, data, systems, people and operating conditions.

A project can begin with an opportunity, a system already in use, a governance need or a research question. The aim is continuity: knowledge that informs the decision, implementation that tests the assumptions, results that guide the next improvement.

The operating cycle

01
Discover
Find and qualify the opportunity: market options, evidence, constraints

→

02
Create
Design and build the solution with the organisation's data, systems and people

→

03
Run
Operate it responsibly, with owners, controls and support

→

04
Improve
Measure performance against the objective and adjust

→

05
Grow
Develop new capability and new sources of value from what was learned

AVAILABLE TODAY

The Good AI Prism research edition and this study; research profiles and comparisons; advisory, assessment, governance and transformation work.

DIRECTION

An AI Business Operating System through which organisations retain intelligence, coordinate their AI estate and develop capability over time. Individual functions are described at their actual stage of availability.

SELECTED REFERENCES & PARTNERS · SINCE 2005

Bosch	Bayer	BMW	Siemens	Mercedes-Benz	Coca-Cola
Deutsche Bank	T-Mobile	Telefónica o2	DHL	EDEKA	AIDA Cruises

Since 2005, the ETHORITY brand and its predecessor businesses have worked with hundreds of brands across industry, finance, consumer goods, energy, travel and telecommunications, including international groups listed in the Fortune 500. Among the 30 companies in the DAX 30 of 2012, 22 appear in ETHORITY's historical client and contact records, about three quarters. Long-running projects included Telefónica o2, Bayer, T-Mobile, EDEKA and AIDA Cruises.

Selected historical references and partners from the group's market-intelligence and advisory heritage, in projects or through former group companies. They describe the group's earlier work.

This background makes the purpose of the study concrete. It combines a view of the market with the questions organisations must answer inside their own operations: what should improve, which options are credible, what evidence is missing, what implementation requires and who owns the outcome.

ABOUT THIS EDITION

Disclosure, quality and corrections

ETHORITY provides advisory, assessment and implementation services related to the topics of this study. Every Prism entry carries its source, the date it was read and the depth of research behind it. Research and drafting were supported by AI tools under the direction of the authors, who retain final editorial responsibility. All web sources were accessed in October 2026. This study is not legal advice.

Editions

This Germany & United States Focus edition is a shorter reading of the same source as the Global Edition, which is the flagship and carries the full global and regional analysis. Regional and sector briefs follow as separate editions. Each edition carries an edition ID in its footer; the matching file fingerprints (SHA-256) will be listed at ethority.ai/ai-prism-study when the edition is released.

HOW THIS EDITION WAS CHECKED

- Review and author approvals confirmed by Sten Franke on 21 October 2026, covering founder and company texts, quotations, portraits and historical reference names.
- Legal dates checked against the operative EU and Colorado texts and the California and India rules.
- Primary releases cited for Brazil (IBGE), Mexico (Banco de México, KPMG) and the Gartner forecast; Eurostat country values checked; vendor-sponsored sources labelled as such.
- Worldwide figures in chapter 8 read at source (WEF, Allianz, PwC, IMF, IEA).
- Prism counts recomputed from the frozen research edition of 20 October 2026 (998 entries, 194 depth-research records, 775 decision values, 216 relationships, 103 re-checked values).

NEXT WAVE

- Wave 2 re-reads the Customer Service pages by product, plan and region and reports every change with its date.
- Provider responses to the correction invitation, dated and published with the next wave.
- Regional and sector briefs and further language editions.

Corrections. Providers and readers can submit corrections and new documentation at ethority.ai/ai-prism-study. A missing published fact is treated as an open question, not as an absent capability.

Note for citation. Quote the Customer Service finding with both criteria: “10 of 107 customer-service AI offerings publicly documented both an EU hosting option and limits on training with customer data: 8 state no training, 2 train only with opt-in consent (Good AI Prism Market Study, Wave 1; provider information reviewed 19–20 October 2026)”. Do not shorten it to training alone. Quote 194 as depth-research records, not as distinct products. Cite as: ETHORITY (2026), Put AI to Work: The Good AI Prism Market Study 2026/27.

REFERENCES

Sources

Every entry links to the page that was read. Sources are also linked where they are used in the text, tables and exhibit notes; provider facts are linked in the appendix. All web sources were accessed in October 2026.

[1] Gartner, Sep 2026 · gartner.com
[2] McKinsey · mckinsey.com
[3] Stanford AI Index 2026 · hai.stanford.edu
[4] Census · census.gov
[5] Cyprus Mail, Eurostat AI use 2025 · cyprus-mail.com
[6] Bitkom · bitkom.org
[7] IDC Europe, Sep 2026 · idc.com
[8] Regulation (EU) 2026/1744 · eur-lex.europa.eu
[9] Bitkom, Sep 2026 · bitkom.org
[10] ifo, Jun 2026 · ifo.de
[11] KfW Fokus 554 · kfw.de
[12] Bitkom/IDC · bitkom.org
[13] DIHK 2026 · dihk.de
[14] Bitkom Cloud Report 2026 · bitkom.org
[15] KPMG, Jun 2026 · kpmg.com
[16] BCG AI Radar 2026 · bcg.com
[17] IW Köln, IW-Trends 2/2025 · iwkoeln.de
[18] Lewis Silkin, Digital Omnibus on AI · lewissilkin.com
[19] Goldman Sachs 10KSB, Mar 2026 · goldmansachs.com
[20] Intuit AI Impact Report, May 2026 · intuit.com
[21] FreshBooks / Wakefield, Sep 2026 · selfemployed.com
[22] Menlo Ventures, Dec 2025 · menlovc.com
[23] McKinsey State of AI 2025 · mckinsey.com
[24] KPMG AI Pulse Q2 2026 · kpmg.com
[25] Deloitte · deloitte.com
[26] Gartner · gartner.com
[27] Eckert Seamans · eckertseamans.com
[28] Colorado General Assembly, SB 26-189 · leg.colorado.gov
[29] CPPA · cppa.ca.gov
[30] Hunton · hunton.com
[31] U.S. Chamber · uschamber.com
[32] AMA, Mar 2026 · ama-assn.org
[33] EBA · eba.europa.eu
[34] RSM, 2026 · rsmus.com
[35] Bain · bain.com
[36] Itron · investors.itron.com
[37] S&P 451 · spglobal.com
[38] Logistics Management, Sep 2026 · supplychain247.com
[39] Capgemini · capgemini.com
[40] Gallup-Walton · nextgeninsights.waltonfamilyfoundation.org
[41] Reuters Institute · reutersinstitute.politics.ox.ac.uk
[42] Wiley · newsroom.wiley.com
[43] Thomson Reuters 2026 · thomsonreuters.com
[44] Gartner · gcom.pdo.aws.gartner.com
[45] NBER · nber.org
[46] The Register, UK Copilot trial · theregister.com
[47] Salesforce · salesforce.com
[48] Erasmus · eur.nl
[49] Ogletree · ogletree.com
[50] LLRX · llrx.com
[51] Help Net Security · helpnetsecurity.com
[52] Stack Overflow · survey.stackoverflow.co
[53] METR · simonwillison.net
[54] Google RCT · arxiv.org
[55] Cleary Gottlieb · clearygottlieb.com
[56] KRDO / Stacker, AI notetaker survey · krdo.com
[57] Eurostat news, 20% of EU enterprises used AI in 2025 · Dec 2025
[58] Bitkom Research, majority of companies use AI · Sep 2026
[59] McKinsey, The state of AI · Aug 2026
[60] California Privacy Protection Agency, ADMT regulations finalised · Sep 2025
[61] Council of the EU, PE-CONS 30/26 (Digital Omnibus on AI) · 2026
[62] Norton Rose Fulbright, Italy Law No. 132/2025 on AI · 2025
[63] World Economic Forum, Global Risks Report 2026 (Global Risks Perception Survey) · 14 Jan 2026
[64] PwC, 29th Annual Global CEO Survey · 19 Jan 2026
[65] Allianz Commercial, Allianz Risk Barometer 2026 · 14 Jan 2026
[66] IMF, Global Financial Stability Report · Oct 2026
[67] IEA, Key Questions on Energy and AI · 2026
[68] IEA, Powering Artificial Intelligence · 2026
[69] heise online, Siemens mit dem besten Markenimage im Internet · 30 Oct 2003
[70] Stanford HAI, AI Index Report 2026, via Burges Salmon · 2026
[71] IBGE, PINTEC Semestral: industrial firms using AI · Sep 2025
[72] Banco de México, Reporte sobre las Economías Regionales, Recuadro 4 · 2026
[73] KPMG México, Global AI Pulse Q2 2026 · Jul 2026
[74] World Bank, World Development Report 2026: The Promise of AI · Aug 2026
[75] World Wide Worx, SA Generative AI Roadmap 2025, via Bizcommunity · Jul 2025

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.

Good AI Prism

The independent map of the AI market: offerings, documented facts and open questions.

Research and studies

Market and decision research that turns evidence into choices organisations can defend.

Advisory and delivery

Readiness, governance and transformation work that puts AI to work, with measurable outcomes.



CONTACT

Sten Franke

Founder, Principal Architect & Lead Creator

sten@ethority.ai · ethority.ai · LinkedIn

Study, exhibits and updates: ethority.ai/ai-prism-study



ETHORITY · SAN FRANCISCO · HAMBURG

© 2026 ETHORITY & STEN.AI LLC · San Francisco HQ · European operations: Hamburg, Germany · ethority.ai · EN-FOCUS-1.10
Put AI to Work: The Good AI Prism Market Study 2026/27 · First edition · Wave 1 of ETHORITY's longitudinal research programme · Edition 1.10, 23 October 2026 · All rights reserved. Exhibits may be quoted with source attribution.

Artwork, sources and reuse

Edition 1.10 · 23 October 2026. Artwork update; the study's evidence cut-offs, findings and edition scope are unchanged. The v3.11 artwork depicts the 17 October reference selection; it is not a count of the study research inventory or the live website catalogue.

Good AI Prism artwork © 2026 ETHORITY & STEN.AI LLC, by Sten Franke & ETHORITY®. CC BY-NC-ND 4.0. Unchanged editorial and press use with credit is additionally permitted. Commercial licences: ETHORITY. Data, software and third-party logos and trademarks are not covered.

CC BY-NC-ND 4.0 · licence terms

The licence above covers the Prism artwork only. Citation does not grant permission to redistribute third-party sources. Study and source rights remain separate. PDF citations are clickable; opening behavior depends on the PDF viewer.