

✦ REPORT · MARKET INTELLIGENCE

AI Industry Trends

What is changing in applied AI, what it means for the organizations buying and building it, and the moves we recommend — from the vantage point of a team shipping production AI every week.

FORMAT	LEVEL	READ TIME	UPDATED
Report	Beginner	12 min	Oct 2026

OVERVIEW

What this report covers

Applied AI has moved from experiments to expectations. Models keep improving, but the bigger shifts are happening around them: agents that take actions, open standards that connect them to tools, evaluation and governance becoming engineering disciplines, and buyers asking for outcomes rather than demos. This report is SplendensLabs' qualitative read of the market, drawn from our client work across healthcare, education, commerce and enterprise operations. It contains no market-sizing figures by design — just the signals we see and what we think you should do about them.

KEY TAKEAWAYS

- The competitive edge has shifted from access to models toward data, workflow integration and evaluation.
- AI is moving from answering questions to completing tasks, which raises the bar on reliability, permissions and oversight.
- Model choice is becoming a routing decision: many teams now use several models, both hosted and open-weight.
- Governance is turning from policy documents into engineering controls that auditors can inspect.
- The winners we see start narrow, measure relentlessly and scale what is proven.

WHO IT'S FOR

- Business and technology leaders
- Product and innovation heads
- Founders building AI-native products
- Anyone briefing a board on AI

HOW TO USE IT

Skim section 01 for the big picture, read the trends relevant to your role in 02–04, and use the recommended moves in 06 to set your next two quarters.

CONTENTS

01 The shift from novelty to infrastructure
02 Trends in models and platforms
03 Trends in how AI is applied
04 Trends in governance and operations
05 What it means for buyers and builders
06 Recommended moves

01 / BIG PICTURE

The shift from novelty to infrastructure

The conversation in boardrooms has changed. The question is no longer whether AI works, but where it reliably creates value and how to run it safely at scale.

THE PILOT ERA	THE PRODUCTION ERA
Chatbots and demos on a single model	Task-completing systems wired into workflows
Success judged by impressions	Success judged by measured outcomes
Innovation budgets and side projects	Line-of-business budgets and owners
Risk handled by restricting access	Risk handled by engineering controls and audits

This shift rewards organizations that treat AI as a capability to be engineered and operated — with data pipelines, evaluation, monitoring and clear ownership — rather than a tool to be licensed and rolled out.

METHOD NOTE

How to read this report

These are qualitative observations from our delivery work and conversations with buyers and builders, not survey results. Treat them as signals to test against your own context, and verify any product-specific details at the time of reading.

02 / TECHNOLOGY

Trends in models and platforms

Models are still improving quickly, but the practical implications are about cost, choice and capability shape rather than headline scores.

SIGNAL

Reasoning on demand

Models that ‘think longer’ when asked have made multi-step analysis and planning practical, with a latency and cost dial you control.

SIGNAL

Capable small and open-weight models

Smaller hosted and open-weight models now handle many routine tasks, making self-hosting and on-device AI realistic options.

SIGNAL

Multimodal by default

Reading documents, images, screens and speech is now standard, unlocking forms, scanned records and voice workflows.

SIGNAL

Routing over loyalty

Teams increasingly route each task to the model that wins their evals, behind a gateway with fallbacks.

Implications for builders

- Design model-agnostic from day one; the best model for a task will change within the life of your product.
- Price-per-token is falling for equivalent capability, but total spend often rises as usage and agentic loops grow — meter cost per outcome.
- Open-weight models are worth evaluating wherever data cannot leave your environment.

✦ 03 / APPLICATIONS

Trends in how AI is applied

The most visible change is in what AI systems are asked to do: less chatting, more doing.

01 **From answers to actions**

Agents now file tickets, update records and draft replies for approval, which makes permissions and audit trails central.

02 **Open standards for tools**

Protocols such as the Model Context Protocol make it easier to connect models to business systems once and reuse that across tools.

03 **AI inside the workflow, not beside it**

Value is shifting to features embedded in CRMs, EHRs, LMSs and storefronts rather than standalone chat windows.

04 **Software delivery itself is changing**

AI coding agents are reshaping how software is written and reviewed, raising the importance of tests, review discipline and architecture.

05 **Vertical depth beats horizontal breadth**

Domain-specific systems — in therapy, education, pet care, retail — outperform generic assistants on real tasks.

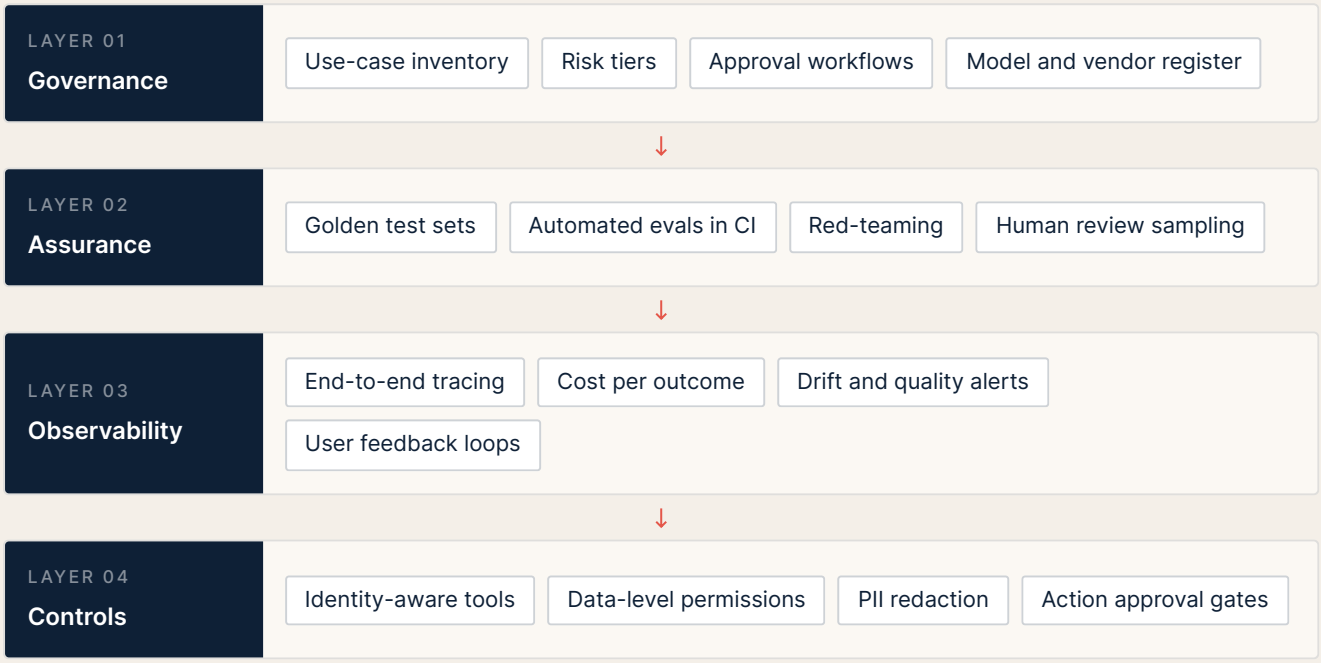
> The interesting question has moved from ‘can the model do it?’ to ‘can we trust the system to do it every day?’
>
> SPLENDENSLABS PERSPECTIVE

04 / OPERATIONS

Trends in governance and operations

As AI touches customers and regulated data, the operational disciplines around it are maturing quickly.

The emerging AI operations stack



Regulatory signals to track

- The EU AI Act's obligations are phasing in, shaping requirements for anyone serving EU users.
- India's Digital Personal Data Protection Act and its rules are raising the bar on consent and data handling.
- Sector regulators in health, finance and education are issuing AI-specific guidance.
- Enterprise buyers are adding AI-specific questions to security and vendor reviews.

05 / IMPLICATIONS

What it means for buyers and builders

The same trends land differently depending on which side of the table you sit.

TREND	FOR BUYERS	FOR BUILDERS
Agents taking actions	Insist on approval gates and audit logs	Design permissions and rollback from day one
Model routing	Avoid contracts that lock you to one model	Build a gateway and per-model evals
Open standards (MCP)	Ask vendors about tool interoperability	Expose systems once, reuse across agents
Governance as engineering	Ask for evidence, not policies	Make evals and traces part of delivery
Vertical depth	Prefer domain-proven partners	Invest in domain data and workflows

What to watch next

01

Agent identity

How agents authenticate and act on behalf of users across systems.

02

Outcome-based pricing

Vendors pricing per resolved task rather than per seat or token.

03

On-device AI

Capable local models changing privacy and latency trade-offs.

06 / RECOMMENDATIONS

Recommended moves

What we would do in the next two quarters if we were in your seat.

STRATEGY

- [ ] Pick two or three high-value workflows and assign a business owner to each

- [ ] Define the outcome metric before building anything
- [ ] Decide what data AI may touch, and where it may be processed

FOUNDATIONS

- [ ] Stand up a model gateway with logging and cost tracking
- [ ] Create golden test sets for each use case
- [ ] Inventory AI use across the organization, including shadow use

DELIVERY

- [ ] Ship a narrow production release, not a broad pilot
- [ ] Add human approval before any irreversible action
- [ ] Review results monthly and scale only what is proven

✦ GET IN TOUCH

Turn trends into a plan

We help leadership teams pick the right AI bets and engineer them into production — RAG systems, agents, automation and AI products. Bring us your priorities and we will map them to a delivery plan.

- AI opportunity and readiness workshop
- Use-case prioritization against your data
- Production roadmap with risks and effort

BOOK AN AI STRATEGY SESSION →

EMAIL
hello@splendenslabs.com

WEB
splendenslabs.com

OFFICE
Sona Incubation Foundation, Salem