

✦ REPORT · MARKET INTELLIGENCE

Agentic AI Adoption Report

Where AI agents are delivering in production, where they are stalling, and the patterns separating durable deployments from impressive demos — SplendensLabs' view from the field.

FORMAT	LEVEL	READ TIME	UPDATED
Report	Intermediate	12 min	Oct 2026

OVERVIEW

What this report covers

Agentic AI — systems that plan, call tools and take actions toward a goal — has become the center of AI roadmaps. Adoption is real but uneven: bounded, well-instrumented agents are quietly running inside operations, while open-ended ‘digital employees’ often struggle to leave the pilot stage. This report sets out what we observe across client engagements and market conversations, the barriers that recur, the architecture patterns that are winning, and the moves we recommend for organizations at each stage of maturity. It is qualitative by design and contains no survey figures.

KEY TAKEAWAYS

- Agents succeed first where the task is bounded, the tools are well defined and success is easy to check.
- The main blockers are reliability, evaluation, permissions and integration — not model intelligence.
- Production designs are converging on workflows with agentic steps, rather than fully autonomous agents.
- Agent identity, scoped permissions and approval gates are becoming standard requirements from security teams.
- Organizations that build evaluation and observability first scale agents faster and with fewer incidents.

WHO IT’S FOR

Business leaders sponsoring automation

CTOs and heads of engineering

Operations and process owners

Risk, security and compliance teams

HOW TO USE IT

Use the maturity model in section 02 to place your organization, read the barriers and patterns in 03–04, and take the recommended moves in 06 into your planning.

CONTENTS

01 Where agents are working today
02 An adoption maturity model
03 What is slowing adoption
04 Architecture patterns that are winning
05 Implications and what to watch
06 Recommended moves

01 / STATE OF PLAY

Where agents are working today

The deployments that last share a profile: a clear task, a small set of reliable tools, a checkable result and a human nearby for exceptions.

WORKING

Customer support triage

Classifying, enriching and routing tickets, drafting replies for agent approval and resolving routine requests end to end.

WORKING

Document-heavy back office

Extracting, validating and reconciling data from invoices, claims, forms and contracts, with exceptions sent to people.

WORKING

Software engineering

Coding agents that write, test and refactor code inside review workflows and CI.

WORKING

Research and analysis

Gathering sources, summarizing and drafting briefings that a person reviews before use.

STRUGGLING

Open-ended autonomy

Agents given broad goals and many tools over long horizons tend to drift, loop or take unexpected actions.

STRUGGLING

Cross-system processes without APIs

Where systems lack clean interfaces, agents become brittle screen-driving scripts.

METHOD NOTE

Signals, not statistics

Observations here come from our delivery work in healthcare, education, commerce and enterprise operations, and from conversations with buyers and builders. Test them against your own context; verify product-specific details at the time of reading.

✦ 02 / MATURITY

An adoption maturity model

Most organizations move through recognizable stages. Skipping stages is the most common cause of stalled agent programs.

STAGE	WHAT IT LOOKS LIKE	TYPICAL RISK	NEXT STEP
1. Assist	Chat assistants answering questions from documents	Over-trust in answers	Add citations and evals
2. Suggest	AI drafts actions; humans execute every one	Review fatigue	Measure acceptance and edit rates
3. Act with approval	Agents execute after a human approves	Approvals become rubber stamps	Auto-approve only proven low-risk cases
4. Act within limits	Agents act alone inside defined scopes and budgets	Scope creep and silent errors	Monitoring, sampling and rollback
5. Orchestrate	Several agents and workflows coordinate across systems	Compounding failures	Strong tracing and ownership per agent

Most production value we see today sits in stages 2–4. Stage 5 is real but rare, and usually built on top of mature stage-4 components rather than designed top-down.

✦ 03 / BARRIERS

What is slowing adoption

The barriers that recur across organizations are mostly engineering and organizational, not model capability.

TECHNICAL

Reliability

Small per-step error rates compound across multi-step tasks, so long chains fail more often than demos suggest.

TECHNICAL

Evaluation

Teams lack test sets for multi-step behavior, so they cannot tell whether a change made things better.

TECHNICAL

Integration

Legacy systems without APIs, inconsistent data and unclear sources of truth slow every project.

RISK

Permissions and identity

Agents acting with broad service accounts alarm security teams, rightly.

COMMERCIAL

Unpredictable cost

Agentic loops make per-task cost variable without budgets and caps.

PEOPLE

Process ownership

Without a business owner who redesigns the process, agents automate yesterday's workflow.

> Agents rarely fail because the model was not smart enough. They fail because nobody defined what done looks like.
>
> SPLENDENSLABS PERSPECTIVE

04 / PATTERNS

Architecture patterns that are winning

Production agent designs are converging. The common thread is constraining the agent where rules exist and freeing it only where judgment is needed.

AUTONOMOUS AGENT	WORKFLOW WITH AGENTIC STEPS
Model decides every step	Code controls the overall process
Many tools available at once	Agents handle bounded judgment steps
Hard to test and explain	Testable, resumable and auditable
Impressive demos, fragile operations	The dominant production pattern we see

A typical production agent path

01 Trigger from system event → 02 Gather context and permissions → 03 Agent proposes action → 04 Policy and risk check → 05 Approve or auto-execute → 06 Log, verify and learn

Building blocks becoming standard

- Durable orchestration with checkpoints so long-running work survives failures.
- Open tool protocols such as MCP, so systems are exposed to agents once and reused.
- Agent-specific identities with least-privilege, scoped and expiring credentials.
- End-to-end tracing of every model call, tool call and decision for audit.

✦ 05 / OUTLOOK

Implications and what to watch

The market is moving quickly. These are the developments we expect to shape buying and building decisions.

WATCH

Agent-to-agent interoperability

Emerging protocols for agents from different vendors to delegate work to each other.

WATCH

Agent identity and authorization

Identity providers and platforms adding first-class support for non-human, delegated agents.

WATCH

Outcome-based commercial models

Vendors pricing per resolved case or completed task instead of per seat.

WATCH

Agent-ready enterprise software

Business applications shipping APIs and tool servers designed for agents, not just people.

What it means for you

IF YOU ARE…	IMPLICATION	PRIORITY
A buyer of agent products	Capabilities vary widely behind similar marketing	Demand evals on your data and audit logs
Building in-house	Integration and evaluation dominate effort	Invest in tool layer and test sets first
In a regulated sector	Approval and traceability are non-negotiable	Design human checkpoints into the process
A software vendor	Customers will expect agent access to your product	Publish well-scoped APIs or tool servers

✦ 06 / RECOMMENDATIONS

Recommended moves

Practical steps for the next two quarters, wherever you are on the maturity model.

01 **Pick a bounded, valuable process**

Choose one workflow with clear inputs, defined tools and an easy-to-check result, and name a business owner.

02 **Build the test set before the agent**

Collect real cases, including edge cases and ones the agent should refuse or escalate.

03 **Start at ‘act with approval’**

Let the agent execute only after a person approves, and measure approval and edit rates.

04 **Earn autonomy with evidence**

Auto-approve narrow categories only once the data shows they are consistently right.

05 **Standardize the platform**

Reuse identity, tool connectors, tracing and evals across the next agents instead of rebuilding them.

Readiness checklist

- ☐ A named owner and success metric for each agent
- ☐ Least-privilege credentials scoped to the agent
- ☐ Cost and step budgets per run
- ☐ Approval gates before irreversible actions
- ☐ Tracing and an evaluation suite in CI

GET IN TOUCH

From pilot to production agents

We design, build and operate agentic workflows that run inside real operations — with permissions, approvals, tracing and evals built in. Bring us one process and we will show you what an agent can safely take on.

- Agent opportunity and readiness assessment
- Bounded-workflow agent prototype
- Production guardrails and observability

BOOK AN AI STRATEGY SESSION →

EMAIL
hello@splendenslabs.com

WEB
splendenslabs.com

OFFICE
Sona Incubation Foundation, Salem