



State of AI Agents 2026



State of AI Agents 2026

Kingy Agent Readiness Score: AI Agents 2026



AI Agent Production Lifecycle

Agent value comes from controlled action, not from longer prompts.

Trigger

User request, event, ticket, lead, alert, scheduled job, or workflow step.

Plan

Goal, constraints, data access, tool route, risk class, and stop conditions.

Act

Tool calls, browser actions, code edits, CRM updates, API calls, and handoffs.

Observe

Logs, traces, intermediate state, costs, model output, and human feedback.

Evaluate

Task success, factuality, policy compliance, side effects, and rollback path.

Improve

Prompt, tool, eval, permission, memory, and workflow updates based on evidence.

Kingy Agent Buyer Framework

A practical lens for separating useful agents from demo theater.

Task Gravity

Does the job repeat often enough and hurt enough to justify an agent?

Tool Risk

What can the agent change, spend, delete, send, schedule, or expose?

Context Quality

Are docs, data, CRM records, files, and policies current enough to trust?

Human Gate

Where must approval, review, escalation, or rollback interrupt automation?

Traceability

Can a human reconstruct why the agent did what it did?

Unit Economics

Does automation save enough time or revenue to survive model/tool costs?

Agent Risk Matrix

The higher the action rights, the stronger the controls need to be.

Read Only

Search, summarize, classify, draft, explain, and retrieve. Lowest operational risk.

Draft + Ask

Prepare replies, tickets, PRs, orders, and campaigns for human approval.

Limited Action

Execute bounded actions with allowlists, spend caps, and audit logs.

Autonomous Action

Run multi-step workflows with rollback, evals, alerts, and incident handling.

Cross-System

Touch CRM, billing, support, code, files, and ads. Integration risk rises fast.

External Impact

Send messages, publish, charge, deploy, delete, or change customer state.