AI AGENTS · 15 NOVEMBER 2026

AI agent scoping checklist.

What to have ready before you ask for a quote

Seven things set the size of an AI agent build, and your app and your data decide most of them. Have these answers ready before you ask for a quote: a price given before anyone has seen the app is a guess.

> **FIRST CHECK** Sometimes the honest answer is no agent. Moving data between systems on a schedule is an automation job, and costs less.

The seven things that set the size

Tick each item once you have it written down. The worksheet on page 3 has room for the answers.

01 Systems it reads and writes

Each system is an integration, with its own sign-in, permissions, data and limits. A documented API is quick to connect. An old system with no API can take longer than the agent.

- [ ] The name of every system the agent reads from.
- [ ] The name of every system it writes to.
- [ ] For each system, whether it has a documented API.
- [ ] Which tools already have an MCP server. The Official MCP Registry lists many, and the agent can use them instead of new code.[1]
- [ ] How people sign in to each system, and what each person may see.

02 Write actions and approvals

An agent that only reads is the cheap half. Once it can change records or move money, each action needs limits, a log and a person who approves it first. OWASP recommends human approval for high-impact actions.[2]

- [ ] Whether the agent only reads, or also changes records or moves money.
- [ ] Each write action, with its limit.
- [ ] Who approves each action before it happens.
- [ ] What each action must log.

03 Data and documents

An assistant is only as good as what it can find. Clean, current documents index quickly. Scans, duplicates, old versions and folders only some people may open all add work.

- [ ] Where the documents live.
- [ ] How clean and current they are: scans, duplicates and old versions.
- [ ] Which folders only some people may open.

04 Evaluation cases

Before launch, the agent should pass real cases from your records with known right answers. Building that set takes time, more so when answers are judgment calls, and it lets you change models later without guessing.

- [ ] Real cases from your records, each with its known right answer.
- [ ] Which answers are judgment calls, and who on your side makes them.
- [ ] The accuracy the agent must reach before launch.

05 Regulated work

The European Commission's summary of the EU AI Act says high-risk AI systems need risk assessment, quality datasets, activity logging, documentation and human oversight before they go on the market. The rules for high-risk uses in certain areas apply from 2 December 2027.[3] Each duty is build work. This is not legal advice.

- [ ] Whether the agent works in a regulated area, checked with your legal contact.
- [ ] Which duties apply: risk assessment, quality datasets, activity logging, documentation, human oversight.

06 The interface

A chat panel in a screen you already have is the smallest interface. A review console, where people approve actions and see each run step by step, is more work, and it is usually what makes a team trust the agent.

- [ ] Where people will use the agent: a chat panel in a screen you have, or a review console.
- [ ] Who approves actions, and who needs to see each run step by step.

07 Users and volume

More users mean more model calls, more support and more ways to fail. An agent your customers use adds tenancy and rate limits. Running costs climb as requests grow.[4]

- [ ] Who uses it: your staff, your customers, or both.
- [ ] How many people, and how many tasks or requests a day.

Ways to keep the cost down

- Start with one workflow that matters, done properly, rather than five sketched.
- Let the agent read before it writes, and put each write behind an approval step.
- Use existing MCP servers for mainstream tools.[1]
- Build the evaluation set first, and scale only once the agent passes it.
- Run the lowest-cost model that passes, and batch the work that can wait.[5]

Scoping worksheet

Fill this in before the first call. Leave a blank where you do not know yet.

App _____ **Date** _____

The one workflow the agent should run _____

SYSTEMS

System	Reads	Writes	Documented API	MCP server	Who may see what
	☐	☐	☐	☐	
	☐	☐	☐	☐	
	☐	☐	☐	☐	
	☐	☐	☐	☐	

WRITE ACTIONS

Action	Limit	Who approves	What is logged

DATA, EVALUATION AND RULES

Where the documents live _____

State ☐ Clean and current ☐ Scans ☐ Duplicates ☐ Old versions ☐ Folders only some people may open

Real cases with known answers _____ **Accuracy it must reach** _____

Who decides the judgment calls _____

Regulated area ☐ No ☐ Yes, which: _____ ☐ Checked with our legal contact

INTERFACE, USERS AND BUDGET

Interface ☐ Chat panel in a screen we have ☐ Review console ☐ Not sure yet

Users ☐ Staff ☐ Customers **How many people** _____

Tasks or requests a day _____ **Budget** _____

Sources, read 15 November 2026

1. Official MCP Registry. registry.modelcontextprotocol.io
2. OWASP Gen AI Security Project: LLM06:2025 Excessive Agency. genai.owasp.org/llmrisk/llm062025-excessive-agency/
3. European Commission: AI Act. digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai
4. InterCode: How Much Does It Cost to Build an AI Agent? Real 2026 Pricing. intercode.com/blog/ai-agent-development-cost
5. Claude API documentation: Pricing. platform.claude.com/docs/en/about-claude/pricing

About Eazetech

Eazetech is an agentic AI studio in Lahore, Pakistan, founded by Babar Ali. We quote each AI build for the app and fix the quote in writing once it is scoped.

The full guide: eazetech.co/pricing/ai-agent-cost