

How an AI agent gets work done.

You set the goal. The system works in steps. You decide what is ready.

YOUR BRIEF
Result + context + allowed actions + approval points + definition of done



THE HARNESS
Software that runs the work around the model.

Model (LLM)
Generates an output or a tool request.

Context
Brief, files, history, and retrieved sources.

Tools
Browser, app actions, search, and code.

Optional subagents: delegate a smaller task and collect its result.

Choose the next step

→

Check access and approvals

→

Run the allowed tool

→

Observe the result

↑

Repeat with the new information

↓

Stop when finished, blocked, or at a time, step, or cost limit.

Good controls: permissions, progress tracking, logs, checks, and stop rules.

YOUR REVIEW + APPROVAL
Check the evidence. Approve before publishing, sending, spending, or changing important data. Verify the final result.

OPTIONAL SUPPORT / ADD WHAT THE TASK NEEDS

RAG
Retrieve sources. Add them to context. Check accuracy.

Skills
Reusable instructions and resources for a type of task.

Plugins
Packages that may bundle skills, tools, and connections.

MCP
A standard for an AI app to connect to tools and data.

INFERENCE

One run of the trained model on the input it receives.

TOKENS

Units the model processes. More context and repeated runs can add cost.

START SMALL: one task, one tool, one clear review point.