

FROM TASKS TO SYSTEMS

An intermediate guide to **AI agents**.

You know how to hand your agent a task. Now learn to build systems that run without you: memory that matters, connected tools, schedules that stick, and projects that manage themselves.

The level-up: stop assigning tasks, start building systems

The beginner move is handing over one task with a finish line. The intermediate move is setting up a workflow that runs without you asking each time.

The shift is small, but everything changes with it: from “do this” to “handle this, ongoing.” A task ends. A system keeps working. Your agent can do both, but systems are where the real leverage lives.

The beginner version

“Check this price now.”

The intermediate version

“Watch this price every day and tell me the moment it drops below \$400.” You asked once. It works forever.

The big idea

Tasks save you minutes. Systems save you hours. This guide is about systems.

Teach it how you work: memory that matters

Tell it the durable stuff once, and stop repeating yourself.

Your agent remembers context across conversations. The intermediate skill is feeding it the things that do not change: how you like your emails (short, no fluff), your budget rules, your projects, the names and details it keeps needing.

When it gets something wrong about you, correct it the way you would a person: “from now on, always use my work email for drafts.” Say it once. It sticks.

Two rules for memory

One: preferences and context go in memory.

Two: secrets never do. Passwords and keys go in the secure vault, never in the chat.

Memory tip

Every few weeks, ask it “what do you remember about me and how I like to work?” Delete what is stale. Memory is a garden, not a hard drive.

Connect your world: inbox, calendar, files

An agent with no connections is a brilliant assistant locked out of your office.

● Calendar

It can schedule, reschedule, and warn you about conflicts before you walk into them.

● Inbox

It can triage, draft replies, and surface only the messages that need your decision.

● Files

It can organize, summarize, and pull the right document the moment you need it.

You grant access once. After that, the approval habit in section 06 decides what it does on its own and what it brings to you first.

The trust ladder

Start with read access everywhere. Let it look before it touches. Expand to write access once it proves it handles the small stuff. Read is cheap. Write is earned.

Put it on a schedule: briefings, checks, watches

Three patterns cover almost every recurring need.

1

The briefing. A digest on your schedule. "Every weekday at 7:30am, send me a 5-bullet briefing: the top news in my industry, my calendar for the day, and anything on my watch list that moved."

2

The check. A sweep that stays quiet unless something real happens. "Every morning, check my inbox. Draft replies to anything routine. Only ping me about the ones that need my decision."

3

The watch. A standing alert with a trigger. "Watch flights to Denver for October. Tell me the moment a round trip drops under \$250."

The rule for all three

Define the trigger, the quiet hours, and what "tell me" means. Vague schedules create noise. Precise schedules create leverage.

Run long projects without losing the plot

Big work dies in long, shapeless conversations. The intermediate move is milestones plus checkpoints.

Before it starts, ask for the plan: "Break this into milestones and show me the plan before you do anything." Approve the plan, then let it run. Then add one rule: "Report back after each milestone with what is done and what is next. Do not wait until the end."

Keep one conversation per project. Everything it learns stays in one place, and you can scroll up any time to see how a decision was made.

When a project stalls

Do not restart from zero. Say "pick up where we left off: the last finished milestone was X, continue from Y." Context is the asset. Protect it.

The approval habit: what to auto-allow vs always check

You do not want to approve everything. You also do not want to approve nothing. Set the line once, in plain words.

Green list: auto-allow

Research, drafts, organizing files, summaries, reminders, price checks. Anything reversible and private.

Red list: always check

Spending money, sending or posting anything, touching accounts, changing settings, deleting anything. Anything hard to undo.

Tell your agent: "Draft anything, send nothing. Research anything, buy nothing. If it is hard to undo, ask me first."

The principle

When in doubt, it asks. That is the deal that makes everything else safe.

When it goes wrong: the 3-minute fix

It will misunderstand you sometimes. The fix is almost always fast.

1

Point at the exact problem. Not “this is wrong” but “the dates are wrong, use October 16 to 18, not September.” Specific corrections teach. Vague ones do not.

2

Give the correction once, then move on. “From now on, prices in USD.” Done. Do not re-explain every time.

3

If the output is vague, ask for shape. “Give me the top 3 with prices and links, not a paragraph.” Agents default to prose. Tell it the format you want.

4

If it assumed wrong, add the constraint it missed. Most bad output is a missing constraint, not a bad agent. “You assumed business travel. This is a family trip with two kids.” Then let it redo.

Do not restart from zero

A correction inside the same conversation keeps everything it already got right.

Five intermediate plays

Copy, paste, customize. Each one is a system, not a task.

1

The Monday briefing. "Every Monday at 8am, send me: the 3 biggest stories in [my field], any price drops on my watch list, and conflicts in my calendar this week."

2

The inbox triage. "Every weekday at 9am, scan my inbox. Draft replies to anything routine. Flag the 3 messages that need my decision, with a one-line summary of each. Send nothing."

3

The project tracker. "Track my [project name] project. Every Friday, ask me for a status update, then reply with: what is done, what is stuck, and what is next."

4

The deal watch. "Watch [product, flight, hotel]. Alert me the moment the price drops below [number] or a clearly better option appears. Include the link."

5

The content multiplier. "Take my [notes / voice memo / rough draft] and turn it into three things: a social post, a short email, and a one-pager. Keep my voice, cut the fluff."

What's next

Pass it on

Know someone still stuck on chatbots? Send them the beginner guide first. New to Muse? Enter the code **DGC4EX** in Settings within 48 hours of joining: you get 1 billion tokens, and so does the person who invited you.

Keep going

This is guide 002.

The advanced guide is next: chaining agents together, building custom workflows, connecting your own tools and data, and running your agent like a small ops team. The free API catalog, the MCP server directory, and the component library are live now.

PLUGMYAGENT.COM // © 2026 HIDDENINWEST