


WHEN ONE AI AGENT BECOMES A TEAM


Supervisors, specialists and handoffs — explained simply.



More agents are not automatically better.

START WITH ONE AGENT

Day 11 built a single agent with a clear job, useful tools, state, and guardrails.



- ✓ Clear goal
- ✓ A few good tools
- ✓ Useful context
- ✓ Progress state
- ✓ Safety limits

Rule: Use multiple agents only when splitting the work solves a real problem.


WHY CAN ONE AGENT STRUGGLE?

Imagine one employee trying to be an accountant, travel planner, support engineer, and security expert — all at once.

Too many instructions

Too many tools

Conflicting permissions




Different expertise

Long, messy context

The answer may be specialization — not a bigger prompt.

WHAT IS A MULTI-AGENT SYSTEM?

Several AI agents with different responsibilities collaborate toward one larger goal.



Shared facts



Clear outputs



Rules & limits

Each agent needs a clear job, inputs, tools, outputs, and limits.

MEET THE TEAM



COORDINATOR

Understands the request, assigns work, and combines results.



SPECIALISTS

Handle narrow tasks using focused tools and knowledge.



SHARED STATE

Carries verified facts, IDs, decisions, and progress.




GUARDRAILS

Control permissions, approvals, budgets, and stopping.

A human still owns high-risk decisions and exceptions.

THINK OF A HOSPITAL


The front desk does not treat every patient. It identifies the need, routes the case, and shares the right details.



Multi-agent systems use the same idea when responsibilities are truly different.

PATTERN 1: ROUTER / TRIAGE

A router identifies the request and sends it to the specialist that should own the conversation.



Best when one specialist should take over and continue directly with the user.

ONE INBOX. THREE DESTINATIONS.

"My card was charged twice."



BILLING
AGENT

"Where is my package?"



DELIVERY
AGENT

"The app crashes after login."

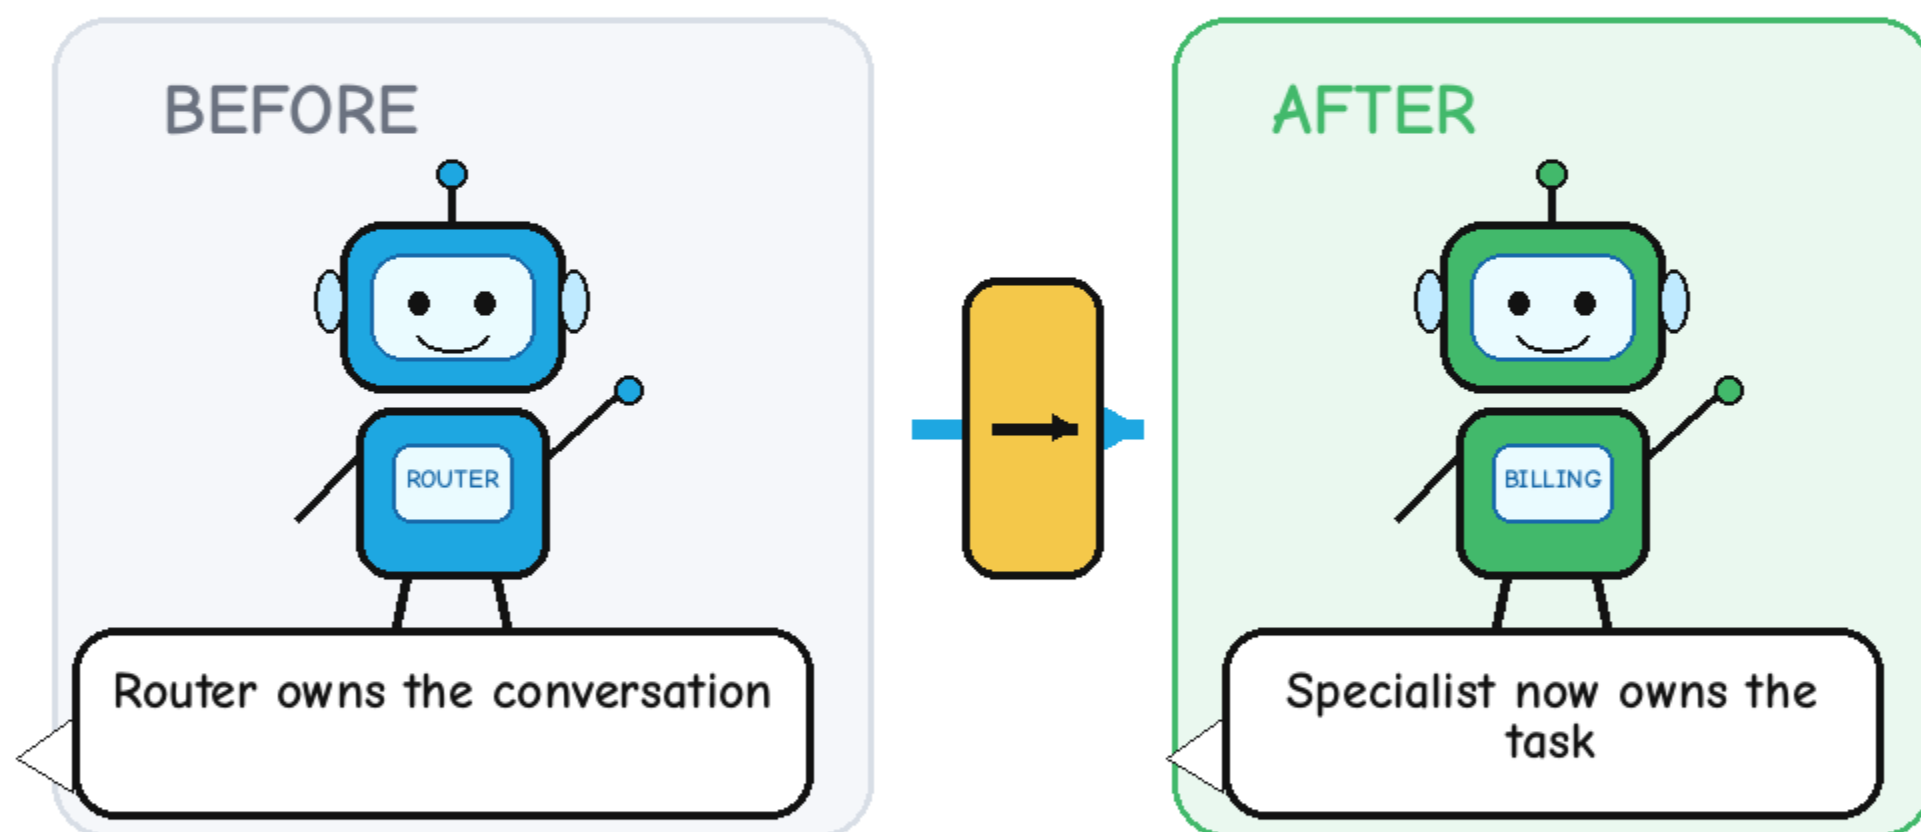


TECHNICAL
AGENT

Same entry point. Different specialist.

WHAT IS A HANDOFF?

A handoff is like transferring a customer call: ownership moves to another agent.



Use a handoff when the specialist should speak directly and continue from there.

A HANDOFF NEEDS A CONTEXT PACK

The next agent should receive the useful facts — not a vague “please help.”



- ✓ User goal
- ✓ Verified identity & permissions
- ✓ Relevant facts and IDs
- ✓ Evidence or tool results
- ✓ Steps already tried
- ✓ Constraints and deadline
- ✓ Expected output

Without this, users repeat themselves and specialists start from zero.

BAD HANDOFF vs GOOD HANDOFF

BAD

"The customer has a problem. Please help."

- No verified facts
- No evidence
- No goal
- No boundaries


GOOD

Order #4821
Duplicate charge detected
Payment IDs A17 / A18
Refund policy verified
Ask before issuing refund

Specific context reduces repeated work and incorrect assumptions.

PATTERN 2: SUPERVISOR / MANAGER


One main agent stays in control while specialists solve bounded subtasks behind the scenes.



The manager delegates, receives results, resolves conflicts, and gives one final answer.

EXAMPLE: PLAN A WEEKEND TRIP

"Plan a 3-day trip to Goa within Rs 30,000."



1
FLIGHT AGENT
Options + prices

2
HOTEL AGENT
Availability + location

3
WEATHER AGENT
Forecast + packing

4
BUDGET AGENT
Total cost

ONE FINAL PLAN

Itinerary + options + total cost + risks

No booking or payment without user confirmation.

HANDOFF OR SPECIALIST-AS-A-TOOL?

HANDOFF

SPECIALIST AS A TOOL

Who owns the task?

Specialist takes over

Manager stays in control

Who talks to user?

Specialist

Manager

Best for

Routing and long conversations

Combining bounded expertise

Example

Refund support


Trip planning

Both are delegation. The key difference is who owns the conversation.

PATTERN 3: PARALLEL SPECIALISTS

Independent subtasks run at the same time, then one agent combines the results.

“Should our company launch Product X?”



Reviews & pain
points

Market
comparison

Price & margin


Legal &
operations

SYNTHESIZED RECOMMENDATION

Faster — but usually more expensive.

PATTERN 4: SEQUENTIAL PIPELINE

Some tasks are naturally completed in stages. Each agent passes a structured result to the next.




Useful for:

marketing campaigns • software release notes •
expert-ready case summaries

Use a pipeline when order matters and each stage has a clear output.

FULL EXAMPLE: ONE REQUEST, TWO PROBLEMS

"My order is missing — and I was charged twice."



ONE ANSWER:

delivery status + duplicate-charge finding + next safe action

Ask before refunding, reshipping, or making another charge.

WHERE MULTI-AGENT TEAMS HELP

1

CUSTOMER SUPPORT

Billing • delivery • technical

2

SOFTWARE

Planner • coder • tester • reviewer

3

RESEARCH

Search • analysis • synthesis

4

HR ONBOARDING

Policy • payroll • IT access

5

FINANCE OPS

Data • reconciliation • compliance

6


HEALTHCARE ADMIN

Scheduling • insurance • records

Use specialists only for distinct expertise, tools, or permissions.

MORE AGENTS CAN MAKE THINGS WORSE

- More model calls mean more cost and latency
- Context can be lost during transfers
- Agents may contradict each other
- Loops and duplicate work can appear
- Debugging and evaluation become harder
- More tools create a larger security surface



A team is better only when the coordination cost is worth it.

DAY 12 TAKEAWAY

Build a team only when specialization or parallelism creates measurable value.

- ✓ Can one agent already do this reliably?
- ✓ Is the task naturally separable?
- ✓ Do specialists need different tools or permissions?
- ✓ Does parallel work save meaningful time?
- ✓ Are handoff inputs and outputs explicit?
- ✓ Is one agent responsible for the final answer?
- ✓ Are budgets, tracing, approvals, and fallback defined?

One clear owner. Focused specialists. Structured handoffs. Strong guardrails.

NEXT UP: Day 13 — Parallel Agents & Deep Research