

How to Use Claude to Level Up in HR

- *A practical guide to building a team of AI employees, not just chatting with a tool.*

01 HR Hasn't Started With Itself

HR is the function every other department looks to when a new way of working shows up. We wrote the change management playbook. We ran the manager trainings.

So it's a little awkward that only 17% of HR departments have a defined AI strategy, and 54% of organizations haven't adopted AI in HR at all. Meanwhile, 92% of CHROs expect adoption to keep climbing anyway.

The function responsible for guiding everyone else through change hasn't finished its own.

This guide skips the theory. You'll learn the one prompting habit that changes everything, why treating Claude as a one-off chat caps what it can do for you, and how to build a small team of named, scoped AI “employees” for recruitment, org design, feedback analysis, and communications.

02 Context Beats Instructions

Most people's first HR prompt looks like this: “Write a job description for a marketing manager.” Claude will produce something generic because you gave it nothing to work with.

Compare that to a prompt with real context: company stage, culture, tone, jurisdiction. Same task, radically different output. This is true of every use case in this guide. Context is the whole game.



The quality of what Claude gives you is a direct function of what you tell it about your company, not how cleverly you phrase the ask.

– on prompting Claude for HR work

Once that habit is automatic, the ceiling isn't the one-off chat. It's what you build next.

03 The Market Already Moved to Digital Employees

A single conversation with Claude forgets everything the moment you close the tab. That's fine for a one-time task, a poor foundation for anything you do repeatedly. The market has already figured out what comes next.

25,000

“digital employees” added by McKinsey alongside its 40,000-person workforce

47%

of leaders using generative AI plan to introduce agentic AI by 2027 – Mercer

7 in 10

businesses plan to integrate AI agents into live workflows this year – Protiviti

One CEO building agentic tools put it plainly: “We don't call them agents, because we think of them more as employees – why we call it ‘onboarding’ and not ‘set up.’”

04 HR's Own Strategy Gap

17%

of HR departments have a defined AI strategy
– SHRM, 2026

54%

of organizations have adopted zero AI in HR, with no
plans to change that

92%

of CHROs expect AI adoption to keep climbing anyway

16%

of AI users qualify as “Frontier Professionals”
– Microsoft Work Trend Index

Individual capability is racing ahead of organizational readiness. The fix isn't a better tool. It's a better practice, and HR has to build it for itself first.

THE CORE IDEA

A tool you configure once and forget is different from an employee **you onboard and hold to a scope.**

Chatting with Claude is useful. An agent-employee with a name, a knowledge base, and a defined scope is what turns one-off help into something your team actually relies on.

05 Build Your First Agent-Employee

A four-step hire, not a technical project.

1

Name & Job Description

A specific mandate, written the way you'd write a real job posting: scope, responsibilities, tone, what success looks like.

2

Knowledge Base

The documents a new hire would need on day one: handbook, past postings, org charts, your brand voice guide.

3

Scope of Authority

What it drafts versus what it decides, written into the instructions from day one, not bolted on later.

4


First-Week Review

Check its output against real examples before you trust it unsupervised. Correct it. Update the instructions.

06 Setting It Up: A Live Walkthrough

What this looks like in practice, inside Claude.ai. No code, no API, about ten minutes.

- 1 Open Projects and create one.**
In the left sidebar, click “Projects,” then “Create project.” Name it after the role, not a generic label.
- 2 Set its instructions.**
Click “Set project instructions” (the pencil icon). Paste in the job description. Every chat in the project inherits it.
- 3 Upload its knowledge base.**
Click “Add content” and upload past postings, scorecards, brand guidelines. Claude reads these automatically.
- 4 Start a new chat from inside the project.**
Not a regular chat. Give it a real task – it already knows the job and has read the files.
- 5 Correct it like a new hire.**
If the first outputs miss, update the project instructions so the same miss doesn't repeat.



Step 1 in Claude.ai: naming a new project.

07 A Ready-to-Use Starting Point

Copy this into a new project's instructions, swap in your own details, and upload two or three past job postings. A working agent-employee in about ten minutes.

SAMPLE PROJECT INSTRUCTIONS — RECRUITMENT COORDINATOR

Role: You are the Recruitment Coordinator for [Company Name], a [size/industry] company. You manage the drafting side of recruitment: job posts, interview guides, candidate communications, and feedback summaries.

Tone: Match the voice of our current job postings and careers page: [e.g., direct, warm, no corporate buzzwords]. Plain language over HR jargon, always.

What you do:


- Draft job posts from a short hiring brief (role, team, must-haves, nice-to-haves).
- Build structured interview guides with a scoring rubric tied to the must-haves.
- Draft candidate emails: acknowledgment, interview invitation, rejection, offer.
- Summarize interviewer feedback into a single comparison view across candidates.

What you never do:

- Never recommend who to hire, reject, or advance. You draft and summarize; a person decides.
- Never send anything externally. Every draft goes to a human for review first.
- Never invent details about a candidate that weren't given to you.

When unsure: ask a clarifying question rather than guessing at role requirements or company specifics.

The same shape (role, tone, what you do, what you never do) is what you'll reuse to stand up each of the other three agents below.



A configured project: instructions saved, ready for its knowledge base.

08 Meet the Team

Four agent-employees, each its own Claude Project with a name, a knowledge base, and a scope.

R The Recruitment Coordinator *manages the full recruitment process*

Drafts job posts, interview guides, candidate communications, and feedback summaries.

Knowledge base: Past postings, scorecard format, employer brand guidelines.

Scope: Drafts everything, decides nothing. No agent says yes or no to a candidate.

O The Organizational Redesign Agent *helps rethink structure*

Models reporting-line options and drafts the rationale memo for a reorg.

Knowledge base: Current organizational chart, headcount, the specific problem being solved.

Scope: Models and drafts, never finalizes. Structural calls stay with leadership.

08 Meet the Team, Continued

The Feedback Analyst *employee engagement*

Turns survey comments and exit-interview notes into themes with real quotes.

Knowledge base: Past survey results, your engagement framework categories.

Scope: Surfaces patterns, doesn't diagnose root cause. A human decides what it means.

The Communications Drafter *employee and organization communications*

Drafts policy announcements, benefits updates, and manager talking points.

Knowledge base: Your comms style guide, past announcements, tone on sensitive topics.

Scope: Drafts every version; compensation, layoffs, or legal exposure get human review.

09 Where to Stop

Build these guardrails into every agent's job description from day one.

1 No Hire/Fire Calls

No agent-employee makes or directly influences a hire, fire, promotion, or compensation decision. It drafts, models, summarizes. A person decides.

2 Legal Review First

Anything with legal exposure, terminations, accommodations, or disciplinary action gets reviewed by qualified employment counsel before it's final.

3 Sensitive Data Stays Out

Health information, disciplinary history, immigration status stay out of knowledge bases unless governance and access controls are built for it.

10 How Big Is Your Team?

A simple self-check, adapted from where most HR functions actually sit today.

Aware

You've used Claude a handful of times for a one-off task. No project, no context saved.

Assisted

You've built at least one Project with real custom instructions and a knowledge base.

Staffed

You have two or more named agent-employees with clear scopes, each doing a job well.

Owned

Your agent-employees are documented, reviewed on a cadence, and used beyond just you.

Most HR functions are sitting at Aware or Assisted right now. The org that gets to Staffed this year will be operating with real capacity most of its peers won't have.

“

HR spent the last decade telling the rest of the business how to manage change. This is the first time the change is happening to us first.

The good news: catching up isn't complicated. It's naming four things, writing them real job descriptions, giving them what they need, and holding them to a scope.

Hire the team you actually need, one agent at a time.

About ScaleHR

ScaleHR helps HR leaders and practitioners build real capability with AI, not just talk about it – through **Catalyst**, our content and influencer marketing program, and community-driven events like **SocialHRCamp**.

Founded by Jeff Waldman after two decades working alongside 200+ HR tech companies, we run on one rule: no buzzwords, no shortcuts, no BS – just work that moves the needle.

Let's Talk → goscalehr.com

If you want help standing up your first agent-employee, or want ScaleHR's take on where your function sits on the maturity ladder above, reach out.

11 Sources

- 1 SHRM, “The State of AI in HR 2026 Report,” survey of ~1,900 HR professionals, published March 31, 2026.
- 2 Microsoft, “2026 Work Trend Index: Agents, Human Agency, and the Opportunity for Every Organization.”
- 3 Gallup, “State of the Global Workplace 2026,” via UNLEASH.
- 4 HR Executive, “Why Every AI Agent Needs a Human Manager and Clear Job Description.”
- 5 DataRobot, “Digital Coworkers: How AI Agents Are Reshaping Enterprise Teams.”