

Building AI agents for the enterprise

Best practices from industry leaders



Contents

Introduction: The thinking divide	4
Upskilling employees	5
Accelerating processes	8
Transforming product development	11
Bringing Claude to the rest of your organization	14
Compounding the AI advantage	19
Resources	21

Foreword

Introduction: The thinking divide

Enterprises are no longer deciding whether to adopt AI, but how fast and at what scale. Our September 2025 [Anthropic Economic Index](#) reports that in the U.S., 40 percent of employees report using AI at work, up from 20 percent in 2023. While these are significant gains, a looming question remains: will AI produce lasting advantages or incremental gains that plateau within a quarter?

The answer depends entirely on scope. Organizations that treat AI as a point solution get point-solution results. A chatbot here, a summarizer there, a pilot that impresses as a demo but never scales (or gets users) beyond the team that built it. This type of adoption initiative can be genuinely beneficial, but it does not change how the organization operates.

Enterprises with successful [AI transformation strategies](#) rethink three things simultaneously: how employees work, how processes run, and what products are possible. This path ensures that AI is not just bolted onto existing workflows but fully embedded into the way your employees work; it guarantees processes that capture and share essential institutional knowledge to benefit the entire org; and it transforms the way customers experience the product.

This is the agentic thinking divide. On one side, organizations deploy AI as a chatbot, a tool that answers questions in isolation and delivers modest results. On the other, they embrace agentic AI, systems that reason through complex tasks, execute multi-step workflows, and apply domain expertise, making it fundamental to how they operate and pull ahead. With increasingly capable models and tooling like skills and plugins, functions beyond engineering, such as legal, marketing, sales, and operations, can now deploy AI that acts on their domain-specific tasks, not just answers questions about them.

In this guide, we explore what leading organizations are doing to drive transformation with AI agents across three key pillars: employees, processes, and products. We also walk through how to implement Claude Cowork and its system of integrated features across your organization, without requiring every team to build from scratch.

Key concepts for business leaders

The language of AI is evolving as fast as the technology itself. This section defines the key terms and concepts referenced throughout this guide, providing a shared vocabulary for evaluating and implementing AI at the enterprise level.

Terms and definitions

Agent

Agents are AI systems that go beyond simple question-and-answer interactions. Where a traditional chatbot responds to a single prompt, an agent can independently plan, make decisions, and use tools across multiple steps to complete a task. Agents determine how to accomplish a goal on their own, selecting the right tools and adjusting their approach based on what they learn along the way.

Chatbot

A chatbot is an AI interface designed for single-turn interactions. A user asks a question, and the chatbot provides a response. Chatbots are effective for straightforward tasks like answering FAQs or retrieving specific information, but they do not plan, take action, or follow through on multi-step tasks the way an agent does.

Claude

Claude is Anthropic's AI assistant. It can read and generate text, analyze documents, write code, and reason through complex problems. Claude serves as the foundation for both chatbot and agentic experiences, and can be extended with tools like MCP, skills, and plugins (see below) to perform specialized work across an organization.

Claude Code

Claude Code is a developer-focused interface that allows engineers to work with Claude directly in their development environment. It can read and edit code,

run commands, search across codebases, and manage complex development tasks, functioning as an AI collaborator embedded in the developer's existing workflow.

Claude Cowork

Claude Cowork is Anthropic's collaborative AI interface that brings the capabilities of Claude Code to everyone. It allows anyone in an organization to work with Claude on complex tasks through a familiar, accessible experience, without requiring coding knowledge or developer tools. Claude Cowork allows knowledge workers to build decks, analyze spreadsheets, compile research, and more tasks across local files, folders, and applications.

MCP

MCP (Model Context Protocol) is a universal standard that allows Claude to connect securely to external services, such as an organization's internal tools and data sources. MCP lets Claude pull from and interact with systems like CRMs, databases, internal platforms, and third-party applications directly, keeping responses grounded in real-time, relevant information.

Plugins

Plugins are pre-built packages that customize Claude for a specific role, team, or organization. They combine the tools, workflows, and knowledge sources a team needs into a single, ready-to-use configuration, so employees get a purpose-built AI assistant without any setup on their end.

Skills

Skills are reusable, predefined workflows that teach Claude how to perform specific tasks consistently. Claude includes a library of built-in skills, and anyone can create custom skills tailored to their specific needs.



Chapter 1

Upskilling employees

Upskilling employees

The earliest AI adopters, including software engineers and financial analysts, were the first to see the value in automating and delegating tasks to LLMs and agents. These productivity gains are now within reach of every knowledge worker. AI is showing up in the tools employees already use—spreadsheets, slide decks, messaging platforms, the desktop itself—and capabilities have moved far beyond autocomplete suggestions and simple drafts.

When they can hand over the repetitive, time-consuming, and information-dense parts of a job to AI, employees can focus on the strategic work that requires their judgement and skillset. A financial analyst stops spending hours pulling data from three different systems and dives into interpreting what the data means. A marketer stops rebuilding the same competitive analysis from scratch every quarter and starts refining the strategy behind it. In this new chapter, expertise does not go away: it goes further.

Tailoring AI to your organization

To work at scale, an organization's AI deployments cannot be generic. A general-purpose AI system that knows nothing about your organization produces general-purpose output that your knowledge workers must edit, deepen, and customize before it is useful. The real gains come when AI reflects how your organization actually works: your standards, your terminology, your tools, your institutional knowledge. That is the difference between AI that drafts a document and AI that drafts a document your team can easily ship.

Consider what this looks like for a sales team. With the right configuration, Claude connects directly to your CRM, your meeting recording platform, and

your prospect research tools. Before a call, an advisor runs a single command and receives a brief that synthesizes the prospect's company data, previous interaction history, open deals, and relevant competitive intelligence. That brief used to take hours of manual research across four or five different applications. In the AI-integrated organization, it takes seconds.

The same pattern applies cross-organizationally across team functions. Your finance team's Claude agents connect to your data warehouse and produce reconciliation reports using your chart of accounts and your reporting standards. Your legal team's Claude-powered system reviews contracts against your specific risk framework and flags deviations from your approved terms. Your marketing team's Claude Project drafts campaign briefs that follow your brand guidelines and pulls performance data from your analytics platform. In each case, the AI is working within your existing institutional context while producing output that satisfies each team's specific needs.

This is why encoding organizational knowledge is powerful. Two companies using the same model will get dramatically different results depending on how much context they give it. The organization that teaches Claude its sales methodology, its compliance requirements, its brand voice, and its reporting standards will see outputs that feel like they came from an experienced colleague.

"Our auto-evaluation capabilities, such as LLM-as-a-judge, have proven many times the superiority of Claude models."

- **Thomas Menard, Head of Agentic Platform and LAB, L'Oreal**

What success looks like

The clearest signal that AI is making employees smarter, not just faster, is when the value is immediate and the setup is minimal. A regional sales manager runs a command before a client call and receives a brief that would have taken an analyst an hour to assemble. A brand manager in a country office asks a question about quarterly performance and gets a visualization with sourced data in seconds, without submitting a ticket to the analytics team. It just works.

Equally important, employees no longer need to be engineers to build and configure the tools that help them do their best work. Specialist knowledge encoded into plugins and shared across the organization means the best practices of one team become the default for every team. The knowledge stops being tribal and starts being institutional.

This is the pattern that separates AI adoption from AI transformation. Adoption puts a tool in front of employees, while transformation changes the baseline of what every employee can accomplish.

L'Oreal

The challenge

L'Oreal's famed skincare, haircare, makeup, and fragrance elixirs are sold in over 150 countries around the globe, and their internal data pipelines were as widespread as their beauty products. Employees who needed data insights had to request custom dashboards for every question. That meant submitting a request, waiting for a specialist to build the query, and hoping the result answered the right question. The painfully slow and technical process was a bottleneck between the teams who knew how to extract data and those who needed it for their everyday work.

The approach

The company built a Claude-based internal AI platform, creating a multi-agent system that transforms natural language questions into data insights and visualizations. The platform was designed so that anyone in the organization, regardless of technical background, could get the information they needed to make better decisions. Claude routes queries to the right data sources and 15 or more specialized agents, then synthesizes results into clear answers with visualizations so employees can now query data directly rather than building custom dashboards for each question.

The results

The platform now serves **44,000 monthly users generating 2.5 million messages per month**, with 15,000 daily active users. Accuracy on conversational analytics improved from 90 percent with previous generative AI approaches to 99.9 percent with Claude. Teams across marketing, product development, and sales now access and analyze information directly, making every employee measurably more effective regardless of technical background. The platform has become core infrastructure for how L'Oreal's workforce operates.

Key metric: Achieved **99.9% accuracy** on conversational analytics



Chapter 2

Accelerating processes

Accelerating processes

AI has the potential to streamline every process in an organization, from generating quarterly planning documents to managing regulatory submissions, legal reviews, and compliance checks. The more complex and information-dense the process, the greater the potential gain. These bottlenecks exist not because teams lack expertise, but because applying that expertise to large volumes of information, with little room for error, takes time.

But process automation is only as impactful as the context behind it. A general-purpose AI model produces general-purpose output. The organizations seeing the biggest gains build their standards, compliance requirements, and institutional knowledge directly into the system. When Claude can reason across an organization's full context, processing times drop from months to minutes while output quality stays high, because the AI is working within the organization's own governed standards and requirements.

The right approach produces compounding returns. As manual processing time and stakeholder coordination shrink, employees redirect that capacity toward strategic and creative work. And as models continue to improve, the processes built on top of them become more efficient without requiring teams to rebuild from scratch.

How accuracy compounds over time

Quality is the most common concern that many enterprises have around adopting AI augmented processes. Speed means nothing if achieving high-quality AI output requires the same number of review-and-improve cycles as the manual processes it's meant to replace.

The organizations achieving the most dramatic process gains with AI avoid this issue entirely by building systems where human expertise feeds back into the

AI's knowledge base. Functionally, these systems propagate curated knowledge: the agents get smarter about domain knowledge over time as subject matter experts review information and make improvements. Each new process starts from the same baseline, and the same review effort is required every time. But in a compounding AI system, every expert review makes every future process better. The customer story below, for example, shows how **Lyft** built subject matter expertise into their customer support AI assistant.

Self-educating, compounding AI systems help the organizations that start earliest build the largest advantage. Every month of accumulated expert approvals, feedback, and refinements makes the next month's output faster and more accurate.

What success looks like

Your organization knows faster processes are working when the people doing the work describe it differently. When, for example, a staff writer who used to spend weeks assembling a clinical study report now explains they now spend most of their work hours on expert review and refinement rather than drafting from scratch. Or when a compliance officer who used to take days to produce a regulatory filing now generates a first draft in minutes and focuses on the judgment calls that require human expertise.

The exact organizational indicators vary role by role, but they all share a predictable pattern. Processes that once required specialized staff become accessible to broader teams, because the AI handles the information-dense assembly work that previously demanded years of domain experience. And the returns compound as more processes come into Claude's scope, because every approved template, every vetted definition, every encoded standard makes the next process faster and more accurate than the last.

The deeper shift is about capacity. When a team that previously spent 80 percent of its time producing documentation can now spend 80 percent of its time on the judgment calls that documentation supports, the nature of the work changes. Senior experts stop being bottlenecked by volume and start applying their expertise where it matters most.

Junior team members gain access to institutional knowledge that previously took years to acquire. The entire team operates at a higher level, because the AI has absorbed the routine complexity that used to consume the majority of their hours.

Lyft

The challenge

Lyft, a global mobility platform operating across six continents and thousands of cities, was facing an inflection point in customer support. The number of customers contacting support was at an all-time high: riders and drivers waited 30 to 40 minutes to reach an agent, and agents were juggling three or four customers at once. The experience felt robotic, with rigid workflows and copy-pasted responses that didn't address customers' actual situations. Agent burnout was rising, and the team lacked the bandwidth to deliver the personal support experience the company wanted.

The approach


After testing multiple AI models on both quantitative performance and brand-voice alignment, Lyft selected Claude to power its customer support AI assistant. The team started with driver support, where onboarding involves region-specific requirements layered on platform-wide policies. From there, Lyft expanded to rider support for charge disputes, ride issues, and other cases previously requiring human investigation. Claude greets customers by name, investigates their specific situation, and resolves issues in seconds. When cases require human judgment, the system routes customers to agents with AI-generated conversation summaries.

The results

Customer support resolution time dropped by over 87%. Decision-making accuracy improved by over 30%. The company reinvested the millions of dollars in savings back into its support agents and new customer programs. Organizing their people and processes around a Claude-powered AI system now lets Lyft:

- Resolve customer issues in seconds rather than the previous 30-to-40-minute wait times
- Free up agents to handle one customer at a time instead of juggling multiple chats
- Fund new programs like Lyft Silver, dedicated one-on-one support for older riders
- Shift support agents from high-burnout volume work to complex issues requiring human care and empathy

Key metric: 87% reduction in customer support resolution time, with **over 30% improvement** in decision-making accuracy.



Chapter 3

Transforming product development

Transforming product development

The enterprises building transformative internal and external products share a pattern: they combine a frontier AI model with proprietary data, existing trust relationships, and deep domain expertise. AI is the enabler, but the defensibility comes from everything around it.

Deploying Claude as an intelligence layer enables new product capabilities that go beyond marginal feature improvements. When your organization combines a frontier AI model with your proprietary data, your existing trust relationships, and your domain expertise, the result is something genuinely difficult for competitors to replicate. Your data, your customer relationships, and your domain knowledge are what make the product defensible. The AI makes it possible to deliver that value in ways that were not feasible before.

The opportunity is not limited to cost reduction. New product capabilities generate net-new revenue and create competitive advantages that compound over time. The organizations that move first establish the integrations, the data flywheels, and the customer habits that make it difficult for later entrants to catch up.

Why the trust boundary determines what ships

There is a critical constraint that determines whether a product idea reaches customers or stays in a proof-of-concept indefinitely: integration within your existing trust and compliance infrastructure.

In regulated industries like financial services and healthcare, data security and regulatory compliance are not optional features. They are prerequisites. Any AI-powered product that operates outside your trust boundary is a product that cannot ship.

In practice, the trust boundary determines which AI integrations are possible and which are non-starters. Financial services firms operate under regulations that specify exactly where client data can be processed, who can access it, and what audit trails must exist.

An AI product that requires client data to leave the firm's security perimeter, even temporarily, triggers compliance reviews that can take months and often end in rejection. The product might be technically superior, but if it cannot operate within the existing trust architecture, it does not get deployed.

The lesson generalizes beyond financial services. In any industry where customers trust you with sensitive data, the trust boundary is the first design constraint for any AI-powered product. Organizations that solve the trust architecture first can move quickly on product development. Organizations that treat it as an afterthought will find that their most promising AI products stall in compliance review indefinitely.

What success looks like

Success metrics for AI-driven product transformation differ from how we evaluate the people and process pillars. With smarter employees, you measure time saved and adoption rates. With faster processes, you measure cycle time compression and quality scores. With transformative products, the question is whether your customers can do something they could not do before.

A clear example: before **Rakuten** adopted **Claude Managed Agents**, our pre-built, configurable agent harness available via the Claude Platform, agentic infrastructure was their constraint. Engineering talent that should have been building differentiated agentic experiences was consumed maintaining the

execution layer underneath. Specialist agents took weeks of custom work to ship, and innovation moved at the speed of the platform team. By offloading that infrastructure entirely, Rakuten now deploys specialist agents within a week, ships major product releases every two weeks instead of once per quarter, and has produced something previously out of reach: a product manager who independently builds FinOps pipelines across multiple public clouds, sets up ambient monitoring agents, and oversees products at a pace the company had never seen.

As Rakuten describes it, the shift is structural: by lowering the cost of coordination and innovation, agents aren't future colleagues or competitors; they're the infrastructure through which the company accelerates everything it builds.

The architectural decision matters just as much as the agent's capability. A product that delivers a genuine capability shift but stalls on infrastructure no one has time to maintain reaches customers too late to matter. Rakuten's decision to offload the execution layer entirely was what let the team focus its investment on the agentic experience itself.

This is where the payoff shifts. Smarter employees and faster processes reduce costs. Transformative products generate revenue. They open new markets, deepen customer relationships, and create switching costs that make the platform more valuable over time. The organizations that invest here earliest will see the largest returns, because the data, integrations, and customer habits they establish today become the foundation for every product they build next.

Rakuten

Rakuten operates over 70 businesses spanning e-commerce, travel, fintech, and communications. As part of its company-wide "AI-nization" strategy, the team saw early that agents would need persistent compute, memory, and storage to move beyond chat-based AI. Their engineers built that infrastructure from scratch—the right call at the time, but one that consumed significant talent that could have gone toward differentiated work.

The deeper goal wasn't running infrastructure. It was democratizing innovation across every business function, so that any employee could delegate outcomes to AI agents.

The approach

Rakuten adopted Claude Managed Agents to offload the execution layer entirely, freeing their engineering talent to focus on the agentic experience itself. The team now deploys specialist agents within a week, covering engineering, product, sales, marketing, and finance, generating apps, proposal decks, and spreadsheets in sandboxed environments.


Agents integrate directly with Slack, Microsoft Teams, and Rakuten's own Kanban-style task system, where employees create and assign tasks to agents from any device. Native mobile and voice support means high-context instructions can be captured on the go. Long-running agents work for hours, allowing teams to delegate goals rather than individual tasks, sharing the desired end state and letting the agent decide what work needs to happen.

Agent memory compounds results over time: agents remember what went wrong in past sessions and avoid repeating those mistakes. Individual learning becomes organizational learning instantly.

The results

Major product releases now ship every two weeks, down from once a quarter. Initial critical errors dropped 97% in pilot, with agent cost and latency down more than 30%, with no loss in output quality. Power users Rakuten calls "Galileo" now contribute across domains far beyond their primary roles: a product manager independently builds FinOps pipelines across multiple public clouds, sets up ambient monitoring agents, and oversees products at a pace the company had never seen before.

The shift Rakuten describes is structural: by lowering the cost of coordination and innovation, agents aren't future colleagues or competitors; they're the infrastructure through which Rakuten accelerates its contribution to society.



Chapter 4

Bringing Claude to the rest of your organization

Bringing Claude to the rest of your organization

What Claude Cowork is (and how it works)

The examples in the previous chapters share a common thread: each organization built a custom platform to give Claude the context it needed to do meaningful work. L'Oreal built an orchestration layer with 15 specialized agents, achieving a 99.9% accuracy on conversational analytics. Novo Nordisk built NovoScribe, reducing clinical study documentation from 10+ weeks to 10 minutes, while maintaining regulatory quality. RBC delivered an agentic solution serving 2,200 advisors managing \$689 billion in assets. Providing advisors more time with each customer. Each investment produced significant returns, but each also required engineering resources, time, and technical expertise to build.

Claude Cowork changes that equation. Claude Cowork is Claude's platform for knowledge workers, giving non-technical professionals the same agent capabilities that enterprises have built on the API, without a custom build. Employees set goals, delegate work, and receive finished deliverables. The output goes beyond drafts and suggestions to actual completed work: Word documents, Excel models, slide decks, analysis, and reports.

The mechanism that makes this practical for enterprises is plugins. Plugins are packages of skills, context, and connectors that give Claude role-specific expertise. A sales team's Claude knows your CRM, your pipeline, and your sales process. A legal team's Claude knows your contract templates, your review workflow, and your risk framework. A finance team's Claude knows your chart of accounts, your reconciliation procedures, and your reporting standards.

Plugins are built once and shared across the organization. When a team encodes its best practices into a plugin, those practices scale automatically to every person who installs it. The institutional knowledge that previously lived in the

heads of your most experienced people becomes available to every new hire, every adjacent team, every office across the organization. Anthropic has open-sourced 11 plugins covering productivity, sales, finance, data, legal, marketing, customer support, product management, enterprise search, biology research, and plugin management itself.

Plugins turn institutional knowledge into organizational infrastructure. Build once, share across every team, and best practices scale automatically.

Built for enterprise scale

Deploying AI across an entire organization raises questions that go well beyond product features. The executives responsible for broad rollout need answers on governance, security, auditability, integration, and workflow continuity. Claude Cowork was designed with these requirements as foundations, not afterthoughts.

Governance and control. Admins configure organization-specific plugin marketplaces, controlling which plugins are available or pre-installed and managing access across teams. This means no unsanctioned adoption, no shadow AI sprawl, and no ambiguity about what capabilities are approved for which roles. Your organization decides exactly what Claude can do, for whom, and under what conditions.

The **plugin marketplace** model addresses a problem that every large organization encounters with AI tools: uncontrolled proliferation. When

individual teams find and install their own AI solutions, the result is dozens of disconnected tools with inconsistent security postures, no central visibility, and no way to enforce organizational standards.

Claude Cowork's marketplace model puts admins in control of the entire catalog. Your IT team decides which plugins are available, which are pre-installed for specific roles, and which require manager approval. When a new plugin is published internally, it goes through your approval workflow before it reaches any employee's machine. This transforms AI governance from a reactive exercise into a proactive one: curating and distributing approved capabilities rather than discovering and shutting down unauthorized tools.

Claude Cowork's admin console also supports role-based access and spend controls, which help organizations manage usage and spend. Usage analytics also provide visibility into how teams are using the tool, including per-user daily activity.

Secure by design. Claude Cowork runs tasks locally, reducing the risk of exposure significantly and making deployment viable without a lengthy security review. For organizations in regulated industries, this architecture addresses the most common objection to enterprise AI adoption before it is raised.

Local execution means exactly what it sounds like: Claude processes your files, your data, and your documents on the employee's machine. Nothing is uploaded to a cloud service for processing.

Auditability. Visibility into what Claude is doing, when, and on whose behalf is available today via OpenTelemetry-compatible observability tools. Your compliance team can see exactly how AI is being used across the organization, which matters both for internal governance and for regulatory reporting. For industries where regulators may ask how AI was involved in a specific decision or document, this audit trail provides the evidence chain that turns a difficult compliance question into a straightforward answer.

Integration with existing tools. Claude works within the systems your teams already use, including CRM platforms, document management systems, and compliance tools, rather than requiring people to adopt a new workflow. The goal

is to make AI feel like a natural extension of the tools your organization already runs on, not a separate system that employees have to context-switch into.

Continuity across workspaces. Work moves fluidly between Claude Cowork and tools like Excel and PowerPoint without losing context. AI output lands where employees already operate. An analyst can start a financial model in Claude Cowork and continue refining it in Excel. A marketer can generate a campaign brief and move directly into slide creation. The work stays connected because the context travels with it.

To learn more about Claude Cowork's enterprise controls, check out our [Claude Cowork Enterprise Admin Guide](#).

Practical deployment framework

The organizations that get the most from Claude Cowork share a deployment pattern that balances ambition with discipline. Four principles guide the most effective rollouts:

- **Start with specificity, not scale.** Give Claude your standards, your tools, your institutional context from the beginning. Employees who receive generic output from a first interaction rarely give the tool a second chance. The organizations in this guide succeeded because they gave Claude enough context to produce output that felt like it came from someone who understood the business.
- **Choose pilots with a measurable finish line.** Each of the three pillars has different success metrics. Smarter employees might be measured by adoption rates and time savings. Faster processes might be measured by cycle time compression and quality scores. Transformative products might be measured by revenue impact and speed to market. Define your success criteria upfront, before the pilot starts, so the results are unambiguous.
- **Build plugins for reuse from the beginning.** The temptation is to build a quick solution for one team and worry about reuse later. Resist that temptation. Plugins built for one team should benefit the entire organization. When you encode tribal knowledge once, every team that installs the plugin gets the

benefit immediately. The marginal cost of sharing a plugin is zero, while the marginal value is enormous.

- **Never underestimate the governance layer.** Admin controls, auditability, and organization-specific marketplaces are prerequisites for broad rollout, not features you add after adoption takes off. The organizations that skip governance early spend more time cleaning up unsanctioned usage than they saved by moving fast.

Phase 1: Setting evaluation and success criteria

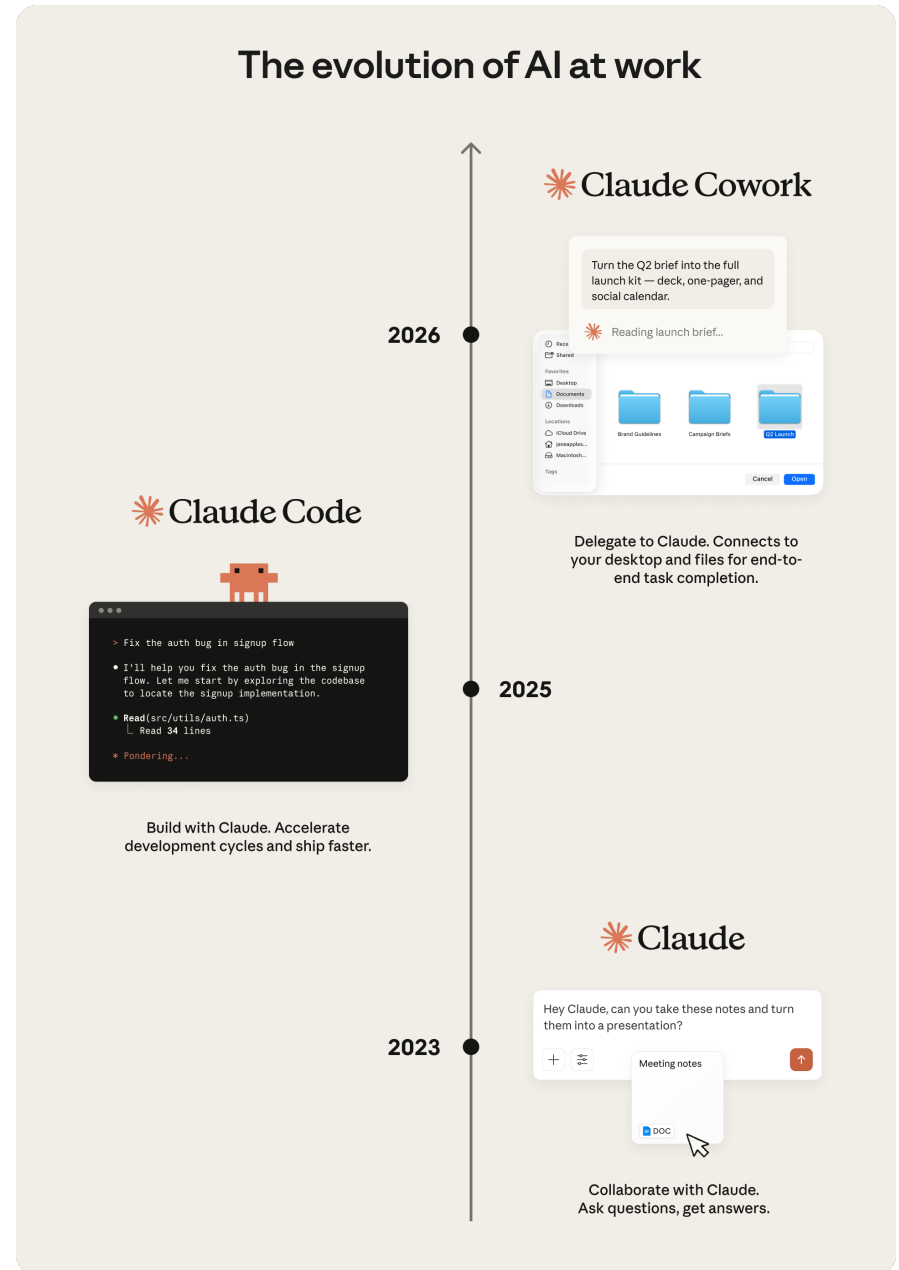
For the first several weeks, focus exclusively on your evaluation and success criteria. Identify two to three teams with clear pain points and measurable workflows. Install the relevant plugins from the [open-source repository](#) or build custom plugins that encode your team's specific standards and processes. Define what success looks like before anyone starts using the tool.

For a sales team, for example, that might be call prep time reduced by 50 percent. For a legal team, it might be a contract review turnaround cut from five days to one. For a documentation team, it might be first-draft quality reaching 80 percent of the final approved version. The specificity of the success criteria matters: vague goals like "improve productivity" produce vague results that are easy to dismiss.

Phase 2: Launching a champion pilot

The second and third months of the initiative are the champion pilot. Two to three teams use Claude Cowork with their configured plugins in production workflows, not sandboxed experiments. Measure adoption weekly. Collect qualitative feedback alongside quantitative metrics, because the moments when employees discover unexpected value are often more informative than time-saved calculations.


The goal of this pilot phase is not perfection but proof of value and a clear understanding of what needs to change before broader rollout. [Getting started with Claude Cowork](#) in the help center provides practical guidance for configuring access and managing this initial deployment.



Phase 3: Scaling impact

Successful proof of concept in hand, months four through six shift to scaling and governance. It's time to deploy the admin marketplace controls, establish plugin review and approval workflows, and begin the rollout to additional teams using the plugins and configurations refined during the pilot.

Each team that comes online benefits from the institutional knowledge already encoded, which means the second wave of adoption moves faster than the first. The third wave moves faster still. This is the compounding dynamic in action: every investment in context, configuration, and governance makes the next deployment cheaper and more effective.



Compounding your
company's AI advantage


Compounding your company's AI advantage

The most common mistake organizations make with AI transformation is waiting until the strategy is comprehensive before taking the first step. The most successful organizations start narrow, learn fast, and expand with conviction. They pick a process, a team, a use case where the pain is obvious and the success criteria are clear. They give Claude the context it needs to do real work, and they measure the results honestly.

Teams moving to centralize AI within their operations now (vs treating agentic AI as an add-on layer) have a compounding head start: trained teams, proven workflows, institutional knowledge encoded in plugins, and a governance infrastructure that can support rapid expansion.

Your organization doesn't need a perfect plan. It only needs a specific starting point, quantifiable success criteria, and the willingness to learn from what happens next.

Contact Anthropic's sales team to scope a pilot tailored to your use cases and define success criteria before you begin. Or **explore Claude Enterprise** for details on security, governance, admin controls, and deployment options that match your organization's requirements.



Resources

Resources

The following resources provide deeper guidance on the tools, strategies, and concepts covered in this ebook.

More on Claude Cowork

- [Introducing Claude Cowork](#)
- [Customize Claude Cowork with plugins](#)
- [Getting started with Claude Cowork \(help center\)](#)
- [Open-source plugin repository](#)

Enterprise AI guides

- [Enterprise AI Transformation Guide](#)
- [2026 State of AI Agents Report](#)
- [2026 Agentic Coding Trends Report](#)



claude.ai