Building an AI-native revenue organization

A revenue leader's guide to running an effective sales organization with Claude

Claude

Table of contents

Foreword: From minutes saved to accelerated revenue growth

For revenue teams, each quarter comes down to hitting or missing a revenue target, and when it's missed, knowing why and how to improve next quarter. To see AI impact at that level, organizations need to adopt and embed it in their processes. Individual reps can save time every day when they hand over call prep to Claude, but when the whole team adopts it, the benefits compound beyond the sum of individual productivity gains.

Rolling out Claude to an organization is different from a traditional SaaS rollout, which typically takes up a big time investment upfront and then settles into maintenance mode. Provisioning Claude seats and training is the first step toward making it a part of everyday work; the returns grow as the organization extends Claude more access and trust. Building an AI-native sales organization is a four-step process, starting with individual productivity gains and time savings as they experiment with AI on their own or in pilot programs, but then advancing through team and department workflows and compounding into KPI or OKR gains.

Cox Communications reported a 7x return in its first-year of AI investment, cutting the cost of validating and enriching sales leads by 86%, and bringing accuracy up from 18% to 97%. At **Cyera**, where 88% of roughly 1,500 employees use Claude weekly, BDRs draft outreach with a skill that's been built using the top performers' best work—a quality improvement as much as a time saving. At Anthropic, **about 80% of our own sales organization** adopted the Claude Sales plugin within months of its release. **One of our sales leaders running a 4,000-account book** estimates that he saves about 90 minutes a day on daily tasks like call prep, and Claude scores every account in his book overnight: work that could otherwise take hundreds of hours across RevOps, FP&A, and marketing.

This guide is for the revenue leader in charge of their department's AI program. We'll walk you through how you can advance your organization's AI maturity, then cover the actual program rollout: your team's setup, connectors, measurement, and common pitfalls to avoid, and wrap up with checklists for getting started.

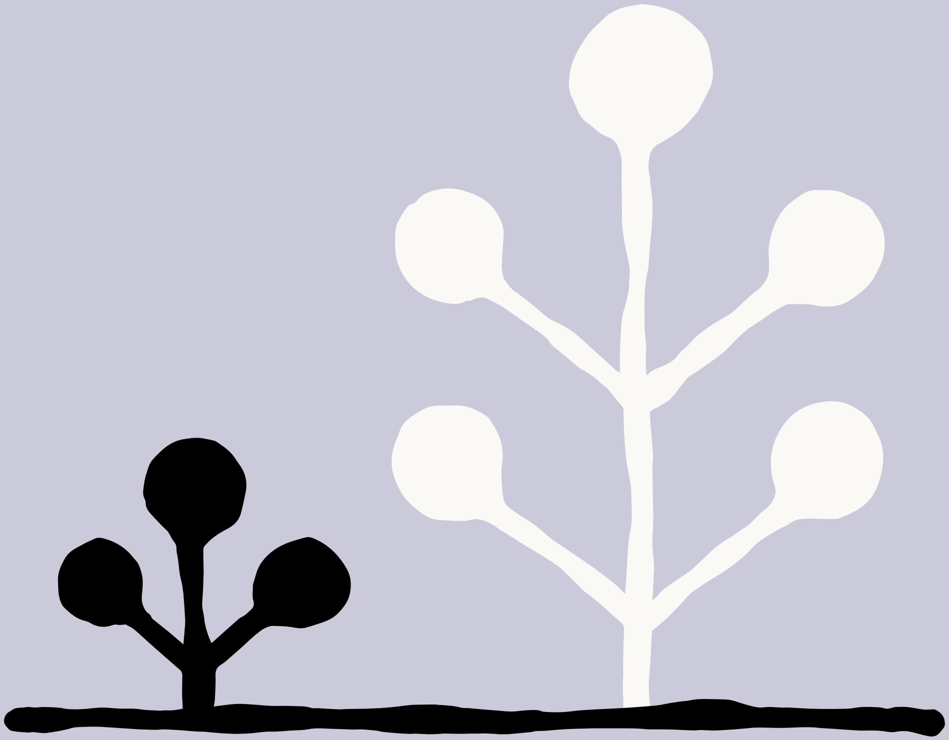

Chapter 1

Advancing your organization's AI maturity

Chapter 1

Advancing your organization's AI maturity

Two revenue organizations can equip the same number of reps with access to Claude and see different returns. What a rep gets from Claude depends on what the organization has built around it: which systems Claude can read and write, how much of the work it is trusted to finish, and whether the team's best methods are written down where everyone can run them. In the deployments we see, impact grows with the complexity of the work Claude can do, and that complexity grows with the access and trust the organization has extended to Claude. The maturity model below makes the progression visible: where your organization stands today, and what they need to do to get to the next level.

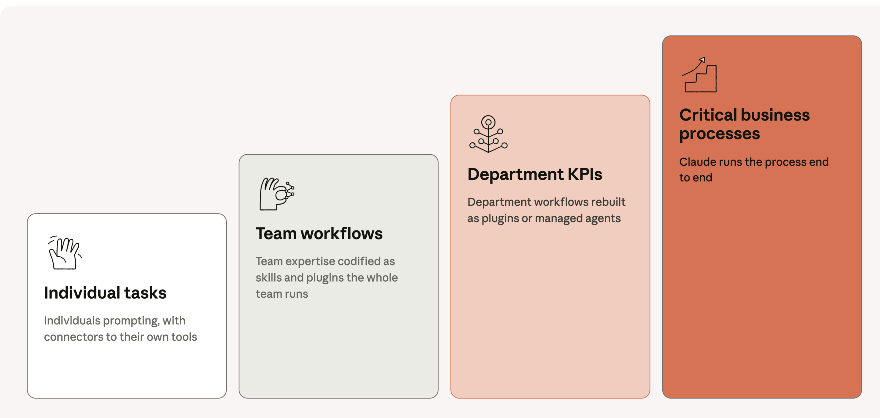

Moving up the maturity ladder

Most organizations start by giving employees a chat assistant, and individual gains show up quickly. Reps can save hours a week on tasks like drafting emails or doing research, but the organizations that move up the that move up the ladder connect systems their teams use, such as Salesforce in Claude (beta) and Anthropic's sales plugin, and extend Claude the trust to take on more complex work. Reps build fluency, from using Claude on single tasks to running several workflows in parallel. Workflows gain access and autonomy, from a rep pasting call notes into a chat to a forecast built from live CRM data. Organizations need to establish governance that supports both: who can publish skills, what data each workflow can reach, and who owns the results. This investment continues throughout the life of the program. More mature deployments result in departments rebuilding workflows as plugins or managed agents pointed at KPIs or OKRs. The goal is to enable Claude to run processes end to end, with human oversight as needed.

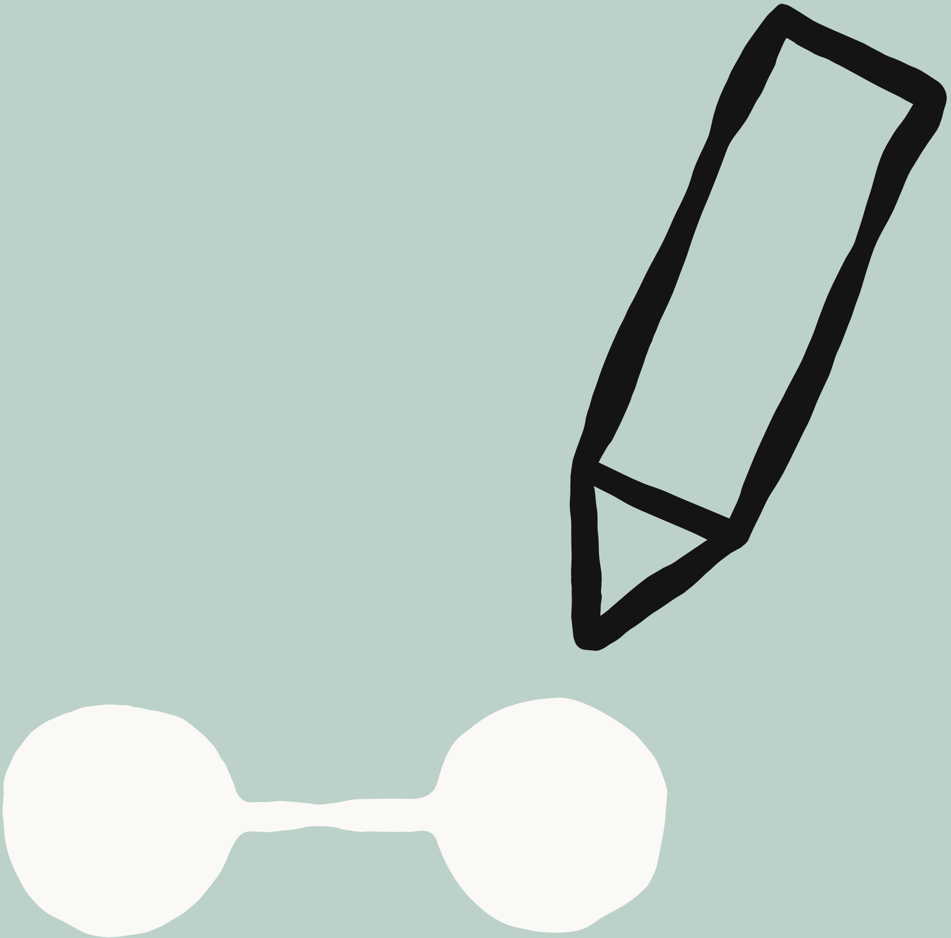

Chapter 2

Setting up for a successful Claude rollout

Chapter 2

Setting up for a successful Claude rollout

Before you get ready to roll out Claude to your sales team, you need to make a few key decisions:

Define owners

RevOps is a strong owner for a rollout because they own the CRM and other systems Claude connects to, as well as reporting on pipeline and forecast, where the results of the program will be measured. At Cyera, the AI team scoped data access and use cases in partnership with RevOps as the system and business owners.

Connect systems

How much your reps can do with Claude depends on how it connects to the systems they use. If your reps need to upload transcript files to Claude after each call, they're spending time on manual tasks that Claude could directly initiate through connectors like Zoom, Granola, or Gong. The more of your systems you connect to Claude, the more work it can finish on its own. We outline some of the connectors sales teams use most commonly in Chapter 3.

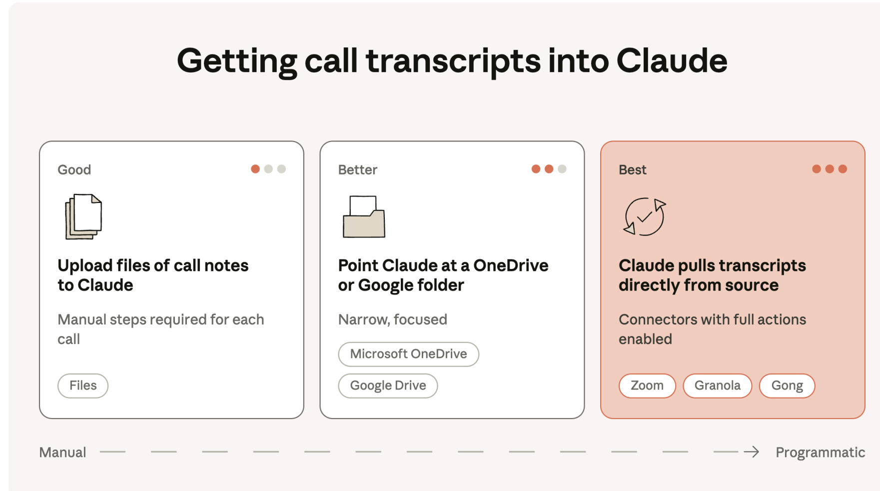

Bring in your IT counterpart

The rollout needs IT for workspace provisioning, SSO, and the admin consents that let connectors reach systems like Salesforce and Microsoft 365. Name an IT owner alongside the program owner, bring them into the security review, and agree on provisioning dates before you commit to a pilot date.

Bring in security

Your security team will need to review Claude before the pilot starts, so start that process early. The review typically covers data boundaries, the connector permission model, auditability, and telemetry export. Our guide, Deploying Claude Across Your Organization, walks you through each of these steps, and Anthropic's security and compliance documentation is at trust.anthropic.com.

Establish success metrics

Pick one activity metric, like call-prep briefs generated or opportunity updates written, and one revenue metric, like pipeline per rep or cycle length, set a baseline ahead of deployment, and commit to a timeline for readouts. (See Chapter 5: Measuring ROI, for our recommendations on setting up baselines.)

Scope spend visibility

Set spend limits by org, group, and user, control access to costlier capabilities by role, and start tracking usage analytics at the very beginning rather than waiting for the first pilot readout. Give champions more headroom than the team default; they run the experiments the rest of the team later inherits as skills.

Select your pilot cohort

Pick two or three teams with motivated leads rather than individual volunteers scattered across the org, and provision plugins at the admin level to keep the pilot from splintering into personal setups.

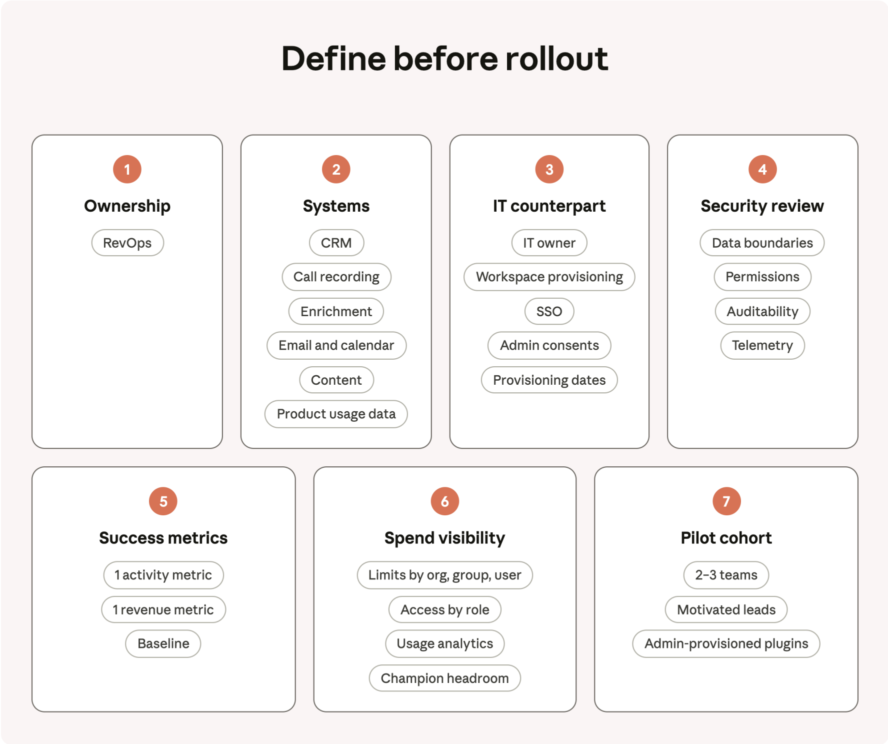

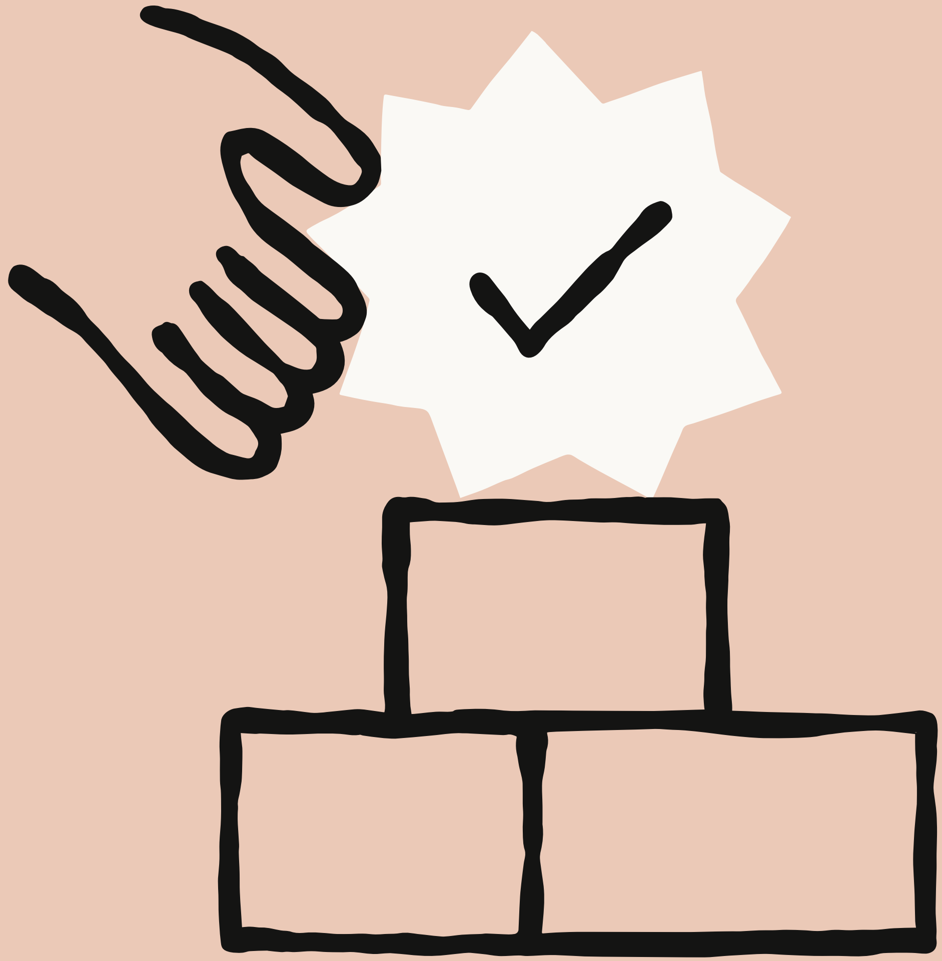

Chapter 3

Your sales teams' Claude setup

Your sales teams' Claude setup

In our conversations with revenue leaders, we hear about the same challenges: half the team barely touches the AI it already has, the best sellers build their own skills and prompts that no one else uses, and buyers have no visibility into actual usage and spend. Account context is scattered across Salesforce, email, Gong recordings, and Slack, and sellers have to reassemble it before every call. That might take thirty minutes of prep for fifteen minutes of conversation, multiplied across a book of hundreds of accounts.

Reps in an AI-native revenue organization can hand all that administrative work to Claude with the commands and skills included in **Anthropic's Sales plugin**.

Commands included in Anthropic's Sales plugin:

Command	Description
`/call-summary`	processes call notes or a transcript, extracts action items, drafts the follow-up, and generates an internal summary
`/forecast`	builds weighted projections from pipeline data
`/pipeline-review`	scores deal health and risk

Skills included in Anthropic's sales plugin:

Skill	Description
`account-research`	Research a company or person — web search for company intel, key contacts, recent news, hiring signals
`call-prep`	Prepare for sales calls — account context, attendee research, suggested agenda, discovery questions
`daily-briefing`	Prioritized daily sales briefing — meetings, pipeline alerts, email priorities, suggested actions
`draft-outreach`	Research-first outreach — research the prospect, then draft personalized email and LinkedIn messages
`competitive-intelligence`	Research competitors — product comparison, pricing intel, recent releases, differentiation matrix, sales talk tracks
`create-an-asset`	Generate custom sales assets — landing pages, decks, one-pagers, workflow demos tailored to your prospect

Initial setup

At initial setup, the Anthropic Sales plugin interviews the rep and commits their name, quota targets, product positioning, and competitor list in a settings file. Although it can be used without connecting Claude to any of your systems (it will work off web search and whatever the rep has pasted in), the plugin only becomes useful when the team connects its tools, so the skills can draw on live data.

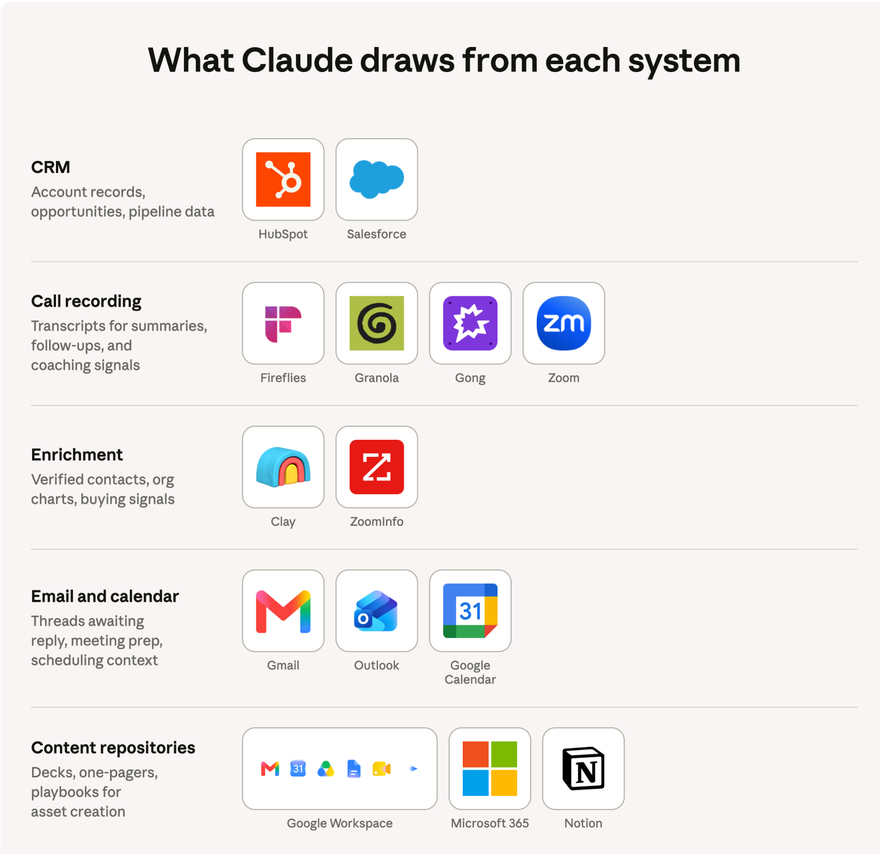

Security

Connectors respect each rep's existing permissions in the underlying systems, so Claude can only access what the rep is allowed to access. If your organization uses Salesforce, admins can provision Salesforce in Claude, which puts a rep's book of business in Claude and runs seller workflows end to end, including renewal prep, QBR decks, pipeline coverage, and meeting follow-up. Since the plugin includes its connectors to Salesforce and Slack, live CRM data is available as soon as the plugin is installed.

Scaling best practices with skills

Your reps can also build their own skills and scale best practices with the broader team. If one of your top performers writes out his renewal-prep routine or forecast format as a skill file, admins can provision it in the team's skill bundle, so new reps start with working methods on day one. Skills are editable files, so if your playbook changes, updating the file will update the skill for the whole team. Learn more about Anthropic's latest security and governance features for skills in our **guide**.

Claude is in the tools your reps already use

The Sales and Salesforce plugins are just two pieces of a wider setup. Claude also works inside the systems your team already uses: **Claude for Excel** for quota and comp models, **Claude for PowerPoint** for proposal and QBR decks, and **Claude Tag** (beta) for questions or tasks your reps communicate in Slack. Here's a rundown of the different Claude surfaces and when to reach for each:

Task	Reach for
A quick question, a rewrite, a brainstorm before a call	**Claude** desktop, web, or mobile
Research, analysis, or a finished document built from your files and connected tools: account briefs, the daily briefing, scheduled runs	**Claude** desktop, web, or mobile
A quota or comp model, or analysis inside a workbook	**Claude for Excel**
A proposal or QBR deck built in the file itself	**Claude for PowerPoint**
A deal question in the team's Slack channel	**Claude Tag**
Code: CRM automations, integrations, data pipelines	**Claude Code**

Using the Sales plugin at Anthropic, our reps can get call prep briefings in about two minutes, where it previously took them thirty minutes of digging through CRM notes, email threads, and call recordings. After each call, Claude drafts follow-up emails and writes the opportunity update straight into the CRM. This replaces the manual data entry that reps used to have to do at the end of the week. As a result, weekly sales forecasts are based on current data, as are any ad hoc checks into the sales pipeline from leadership or other teams.

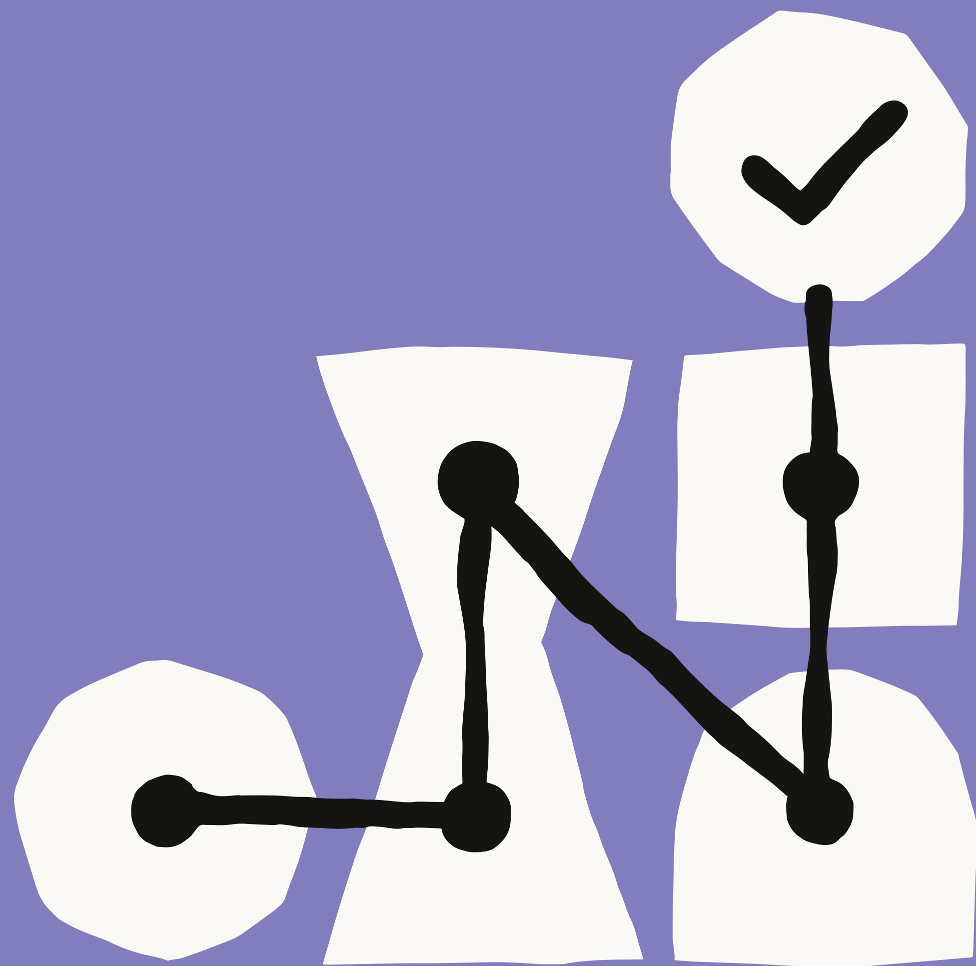

Chapter 4

Your Claude rollout plan

Chapter 4

Your Claude rollout plan

A sales team can go from first seats to a fully provisioned team running shared skills in three phases, progressing through a short setup, piloting on two or three teams, and a scaling phase that extends to the rest of the organization. Here's the plan at a glance:

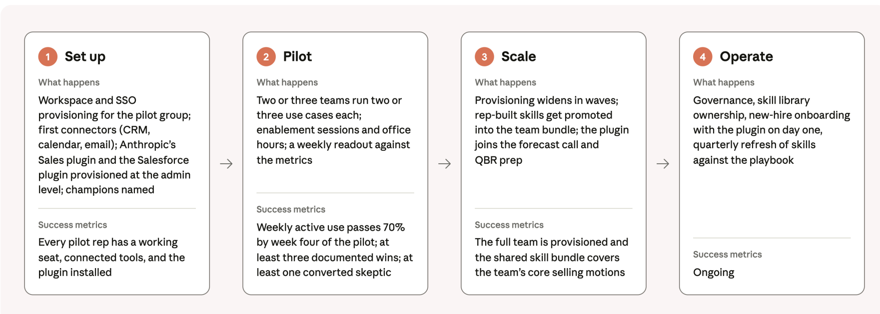

During the setup phase, two groups work in parallel: IT admins handle workspace provisioning and SSO, while RevOps configures the first connectors and provisions plugins at the admin level, using access decisions from security review.

Simultaneously, you need to select champions to run the pilot. Consult with your sales managers here, they know which reps already experiment with AI and which teams have motivated leads. Plan for two or three champions per department, or one for every 25 to 50 users, and give them dedicated hours for this work. During the pilot, they will demo on live accounts, staff the office hours, and write the first skills. The number of champion-created skills that pilot teams use regularly is the best early signal of how scaling will go.

Best practices from the field:

Cyera ran a full-day kickoff livestream with 20 department-specific sessions, followed by twice-weekly office hours to support its rollout.

Cox seeded champions across its teams and scaled through train-the-trainer.

How well your organization will adopt AI depends on their first hands-on experience with it. To make adoption easier, provision a starter pack of plugins and skills that are specific to how your teams work. If you include Anthropic's sales plugin, preload the settings file before the first login, so that every rep has a command that works on day one. For the Salesforce plugin, an admin-enabled setup skill does all this work on the first run: it identifies the rep's tools and connectors, and configures the plugin for them.

During the scale phase, the skills that reps built during the pilot will move into the shared team bundle, and the rest of the team gets provisioned in waves (25 seats, then 150, then everyone is a pattern we see often). This is also when teams will start using Claude's outputs in their standing meetings and shared communication channels: reps bring /forecast projections to the team's weekly forecast call, managers build QBR decks in Claude in PowerPoint, and so on.

Success stories from the field:

After a brief pilot, **Cyera's AI team rolled out Claude to the entire company in 17 days**, with 40 tools connected. Before rollout, the team ran data-mapping exercises: which systems Claude would connect to, which contained production data, and how Cyera would maintain observability. Close partnership with the data and RevOps teams made sure the tool would provide actual value to a wide range of employees.

Use cases by role

Each pilot team runs two or three use cases to start; the rest can wait for the scale phase. Here are the ones we see revenue teams run most, and who gets the most from each:

Daily prep

The daily-briefing readout: meetings with prep, pipeline alerts, email priorities

AEs · SDRs

Call prep and follow-up

A pre-call brief from CRM, email, and call history; after the call, the follow-up email and the CRM update

AEs

Account research briefs

A researched brief on any company or contact, from web search plus enrichment and CRM data

AEs · SDRs

Outreach at scale

Researched, personalized email and LinkedIn drafts from the draft-outreach skill

SDRs · BDRs

Account scoring

Every account in a book scored overnight against your criteria

RevOps

Pipeline review and weighted forecast

/pipeline-review deal-health flags and /forecast weighted projections ahead of the forecast call

RevOps · Sales leadership

Sales reports and QBR readouts

QBR decks and reports built in Claude in PowerPoint from CRM and usage data

Managers · RevOps

Proposal decks and custom assets

Prospect-specific decks, one-pagers, and demo assets from create-an-asset

AEs

The team playbook as a skill

Top performers' methods written as shared skills the whole team runs

Sales leadership · RevOps

The deal desk

Deal questions answered by Claude Tag in the team's Slack channel

Whole team

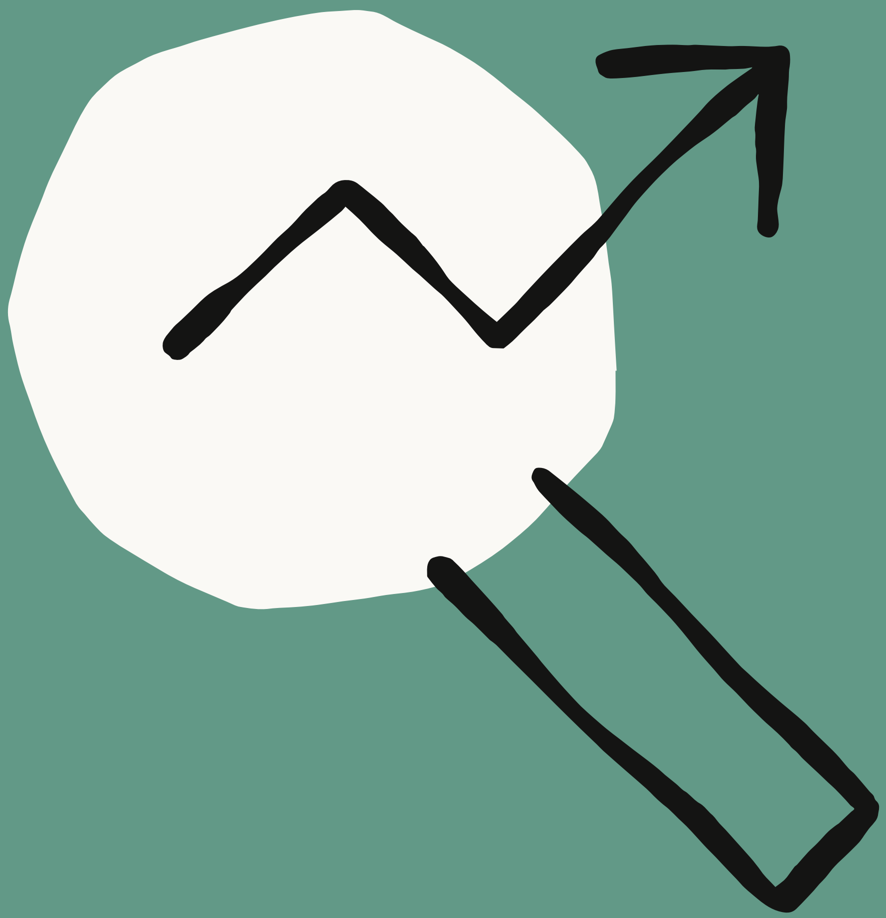

Chapter 5

Measuring ROI

Chapter 5

Measuring ROI

Few sales organizations have defined ROI metrics for AI, and the readout executives most often ask for is return by function. Because you committed to tracking one activity and one revenue metric during the pilot, you can provide clear readouts to your leadership that will help flag any blockers and inefficiencies, while celebrating wins to drive further adoption.

Claude's returns typically show up in four stages: usage moves first, then output, then outcomes that you can track in the CRM, and cost.

What to measure

	Example metrics	Where to read them
1 **Adoption**	Share of seats active weekly, connectors authorized per rep, skill runs per rep per week	Usage analytics in the admin console
2 **Activity**	Call-prep briefs generated, opportunity updates written, proposals drafted — each paired with a quality check, like forecast accuracy or how much of a draft survives the rep's edit	Usage analytics joined with spot checks
3 **Outcomes**	Pipeline per rep, accounts covered per rep, cycle length, win rate, new-rep ramp time	The CRM, comparing adopters with non-adopters
4 **Cost**	Spend per rep per day, cost per unit of work (per brief, per opportunity update)	Usage analytics joined with the activity counts

Sort the returns into three groups: efficiency, expansion, and new capabilities.

- **Efficiency** is the same work done faster: **at Workato, deal prep fell from three or four hours to 45 minutes**.
- **Expansion** means more output from the same team, like more accounts covered and more pipeline per rep.
- **New capabilities** unlock work that didn't happen at all before, like reaching out to the long tail of accounts no rep had time to touch.

During the pilot stage, your pilot groups are selling with Claude while the rest of your teams are not, and your CRM tracks the same metrics for both. Rather than comparing a team's outputs before and after using Claude, compare these two groups side by side to get a sense of the pilot's ROI across several metrics, including pipeline per rep, cycle length, and win rate for the same quarter. Both groups sell into the same market in the same season, so a difference is easier to credit to the program than a change from one quarter to the next.

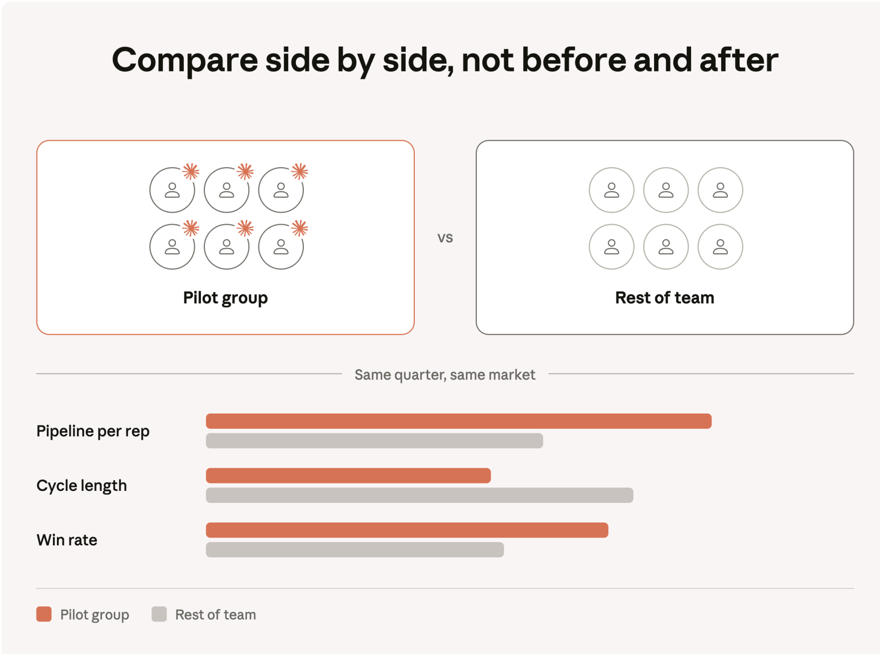

As you wrap up your pilot, three signals justify continuing on to the scaling phase: reps are still producing even after the novelty has worn off, the quality checks are holding, and the pilot teams are outperforming on the outcome metric you picked, even if the gap is still small.

After a full quarter of usage, the program should have enough history for you to look at cost vs results: what did each rep spend with Claude per day, next to pipeline per rep, cycle length, and win rate. Take those stats into the budget conversations for next year and lead with expansion and new capabilities. If you build the case for AI on hours saved alone, the value tops out at the cost of those hours, because the CFO can only value “time freed up” at what the team is paid for that time.

The first time you pull a spend report, you may focus on the reps with the highest spend. That alone isn’t a reliable metric. Instead, read each rep’s spend in the context of what that rep produced: briefs, opportunity updates, proposals. A champion near the top of the table who is also producing every day is the program working as designed; the extra headroom champions got at setup exists for exactly that. When a rep spends heavily and produces little, start with coaching: have a champion walk them through the commands, skills, and connectors they should be using. Lower their cap only if coaching doesn’t change their output.

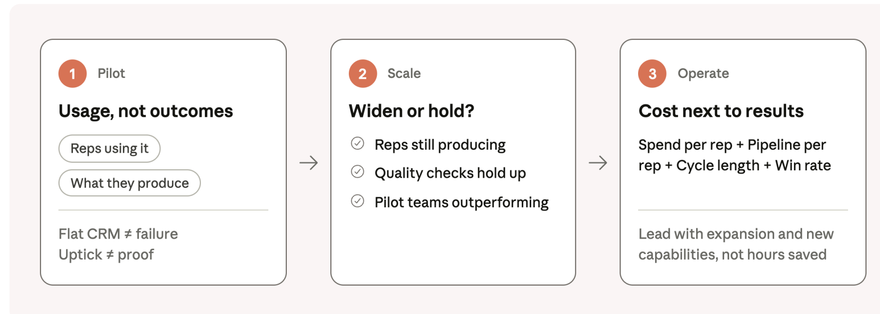

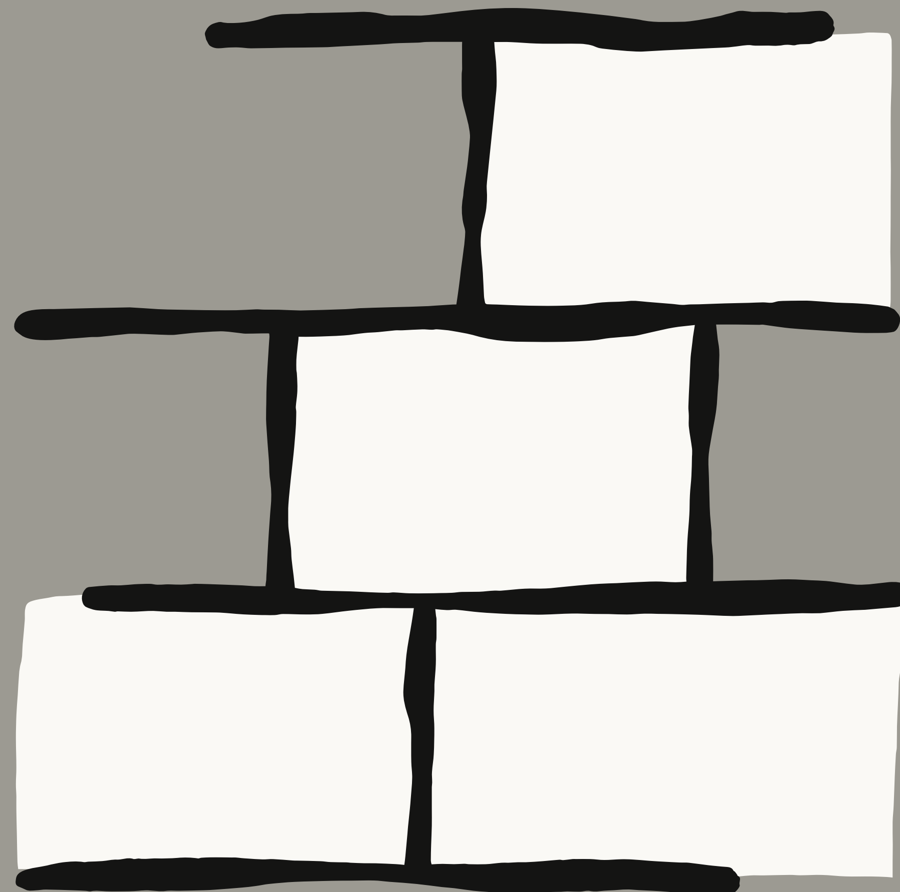

Chapter 6

Common barriers to adoption

Common barriers to adoption

You can run a successful pilot and still stall during the org-wide rollout phase. We’ve seen several patterns that you can plan for during setup:

Running the pilot without an end date

The pilot goes well, so the team keeps it going. This can happen if nobody has set a date for making the decision to move to the scale phase, or named the person who makes that decision. The sponsoring executive will move on to next quarter's priorities, and the rest of the organization will never get provisioned. A year later the program is still 25 seats, and its gains are too small to reach the numbers the CFO reads.

Avoid this pitfall at setup: set a date for a scale decision on the sponsoring executive's calendar, name the person who makes that decision, and agree on what a “yes” requires before you start rolling out to the entire organization.

Scaling seats without scaling champions

Two or three champions can support a 25-seat pilot. The mistake is widening provisioning to 150 seats while the champion count stays at three: office hours are hard to accommodate for this many people, and the reps provisioned in later waves could get a bad first week with nobody around to help them ramp up.

Hold the ratio through each expansion wave, one champion for every 25 to 50 users. Before each wave gets provisioned, the program owner should ask the incoming teams' managers to name their champions, so the current champions can train them.

Ignoring spend until the first invoice

With usage-based billing, monthly costs rise as more reps get provisioned and run more work. A common mistake here is not setting spend controls: if nobody sets limits or reads usage analytics at a regular cadence, the budget owner will learn the cost from the monthly invoice. When that number is higher than planned, the common response is to restrict access or remove users, which interrupts adoption in the middle of the scale phase.

Set spend limits by org, group, and user before the pilot, and read the usage analytics weekly. A budget owner who watches the spend curve build can budget for it before the invoice arrives.

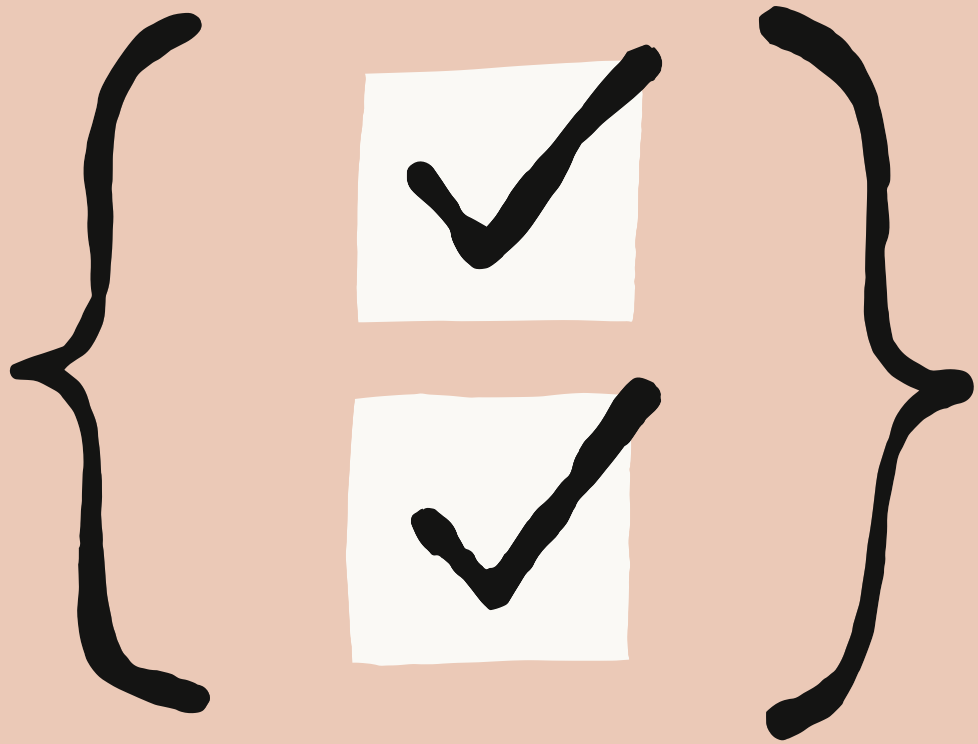

Chapter 7

Getting started

Chapter 7

Getting started

To get started, follow the steps laid out below. If your team is already working with Claude, skip to the steps that best match your team's progress today.

Phase	Action	Done when
1 Setup	Complete the setup checklist	☐ Every pilot rep has a working seat with the plugin and has run a first command

	Action	Done when
1	**Name the program owner and the sponsoring executive** CRO	☐ Both are named and announced to the org
2	**Name the IT owner and schedule provisioning** Program owner, with IT leadership	☐ An IT owner is named and provisioning dates are on IT's calendar
3	**Book the security review** Program owner, with the security team	☐ The review is scheduled and its scope agreed
4	**Pick one activity metric and one revenue metric** Program owner and RevOps	☐ Both metrics are written down, with reading dates at day 30, 60, and 90 with reading dates set for each phase checkpoint
5	**Put the scale decision on the calendar** Program owner	☐ The sponsoring executive has the date, and the decision maker is named
6	**Set spend limits and usage analytics** RevOps or IT admin	☐ Limits set by org, group, and user; analytics visible to the budget owner
7	**Choose two or three pilot teams and their champions** Program owner, with sales managers	☐ Champions are named, with dedicated hours
8	**Provision the pilot group** IT and RevOps	☐ Every pilot rep has a working seat, SSO, first connectors, Anthropic's Sales plugin and Salesforce in Claude installed at the admin level
9	**Run the first live demos** Champions	☐ Each pilot rep has run their first command on a real account, starting with account-research

Phase	Action	Done when
2 Pilot	Champions demo on live accounts, staff office hours, and write the first skills; each pilot team runs its two or three use cases	☐ Weekly readouts are running and the first champion-written skills are in use
3 During pilot	Read adoption and activity: seats active weekly, briefs and updates produced, each with its quality check	☐ Usage is moving; a flat CRM reading is expected at this point
4 Pilot	Keep enablement running; track which champion skills teams use regularly; document wins and converted skeptics	☐ Weekly active use passes 70% by week four; at least three documented wins; at least one converted skeptic
5 End of pilot	The named decision maker makes the scale call against the criteria agreed at setup	☐ A yes widens provisioning; a no ends or reshapes the pilot, with the reasons written down
6 Scale to organization	Provision in waves; hold the champion ratio (one per 25 to 50 users); managers nominate, current champions train; promote rep-built skills into the team bundle; read spend weekly	☐ The full team is provisioned and the shared bundle covers the team's core selling motions
7 Rollout completed	Put cost beside results: spend per rep per day next to pipeline per rep, cycle length, and win rate	☐ The budget case for next year is written, led by expansion and new capabilities
8 Ongoing	Assign governance and skill library ownership; onboard new hires with the plugin on day one; refresh skills quarterly against the playbook	☐ Ongoing

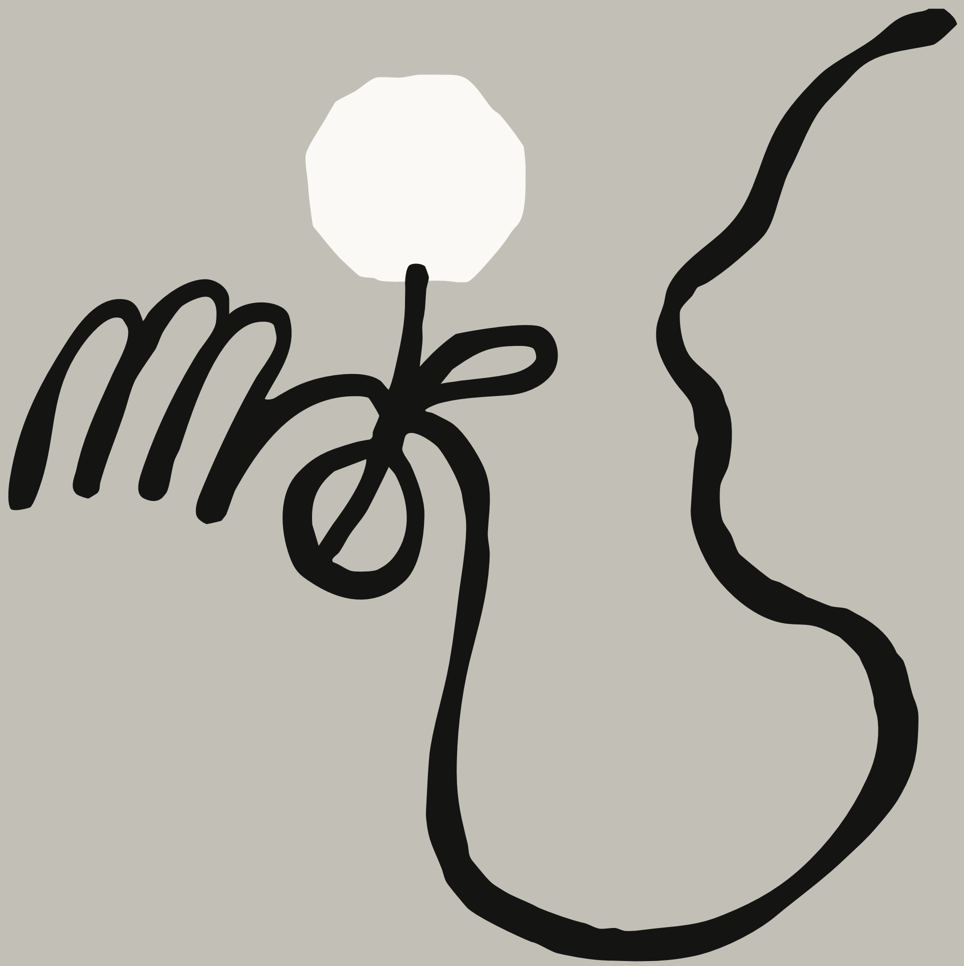

Chapter 8

FAQs

Chapter 8

FAQs

Will Anthropic's Sales plugin work with our CRM and call data?

Anthropic's Sales plugin connects to tools like HubSpot, Clay, ZoomInfo, Fireflies, and Microsoft 365. Salesforce has a dedicated plugin, Salesforce in Claude, which includes the Salesforce connection itself, with live CRM data available as soon as the plugin is installed. Call recordings connect the same way, so transcripts feed call summaries and coaching signals. Connectors respect each rep's existing permissions in those systems, and the connector directory has the current list.

Can I trust the forecast numbers and proposal claims Claude produces?

Claude builds forecasts from the pipeline data in your CRM, so every number traces to a record your team can open. Keep the quality checks from your measurement setup in place: forecast accuracy tracked over time, and a rep review before anything customer-facing ships.

How is this different from the AI already bundled in our software suite?

An assistant embedded in one application helps inside that application. Claude works across the systems a deal touches, the CRM, call recordings, email, and your content repositories, and produces finished work from them: briefs, opportunity updates, proposal decks. Your team's own methods become shared skills, so the setup improves as your top performers write down what works.

How do we forecast and cap spend?

Set spend limits by org, group, and user before the pilot, and control access to costlier capabilities by role. Usage analytics show spend per user and per skill from the first week, so the budget owner sees the curve building well before the invoice. Give champions more headroom than the team default; their experiments become the team's skills.

What are teams seeing beyond time saved?

Hours saved are the floor. The larger returns are expansion, more accounts covered and more pipeline per rep from the same headcount, and new capabilities, work that didn't happen before, like every account in a book scored overnight. For specific examples, see:

- Cyera scales agentic AI across 1,500 employees with Claude Enterprise
- Cox Communications drives a 7x return on AI across its B2B funnel with Claude and Accenture
- Workato connects enterprise AI agents to business systems with Claude
- How an Anthropic sales leader uses Claude Cowork to run a 4,000-account book
- and more in the Anthropic customer story library

Will our CIO and CISO approve it?

The security review your team runs before the pilot. Chapter 2, Setting up for a successful Claude rollout, covers what they will ask: data boundaries, the connector permission model, auditability, and telemetry export. Claude only reaches what each rep is already allowed to reach, because connectors inherit the permissions of the underlying systems. Our guide, **Deploying Claude Across Your Organization**, walks through the review, and Anthropic's security and compliance documentation is at **trust.anthropic.com**.

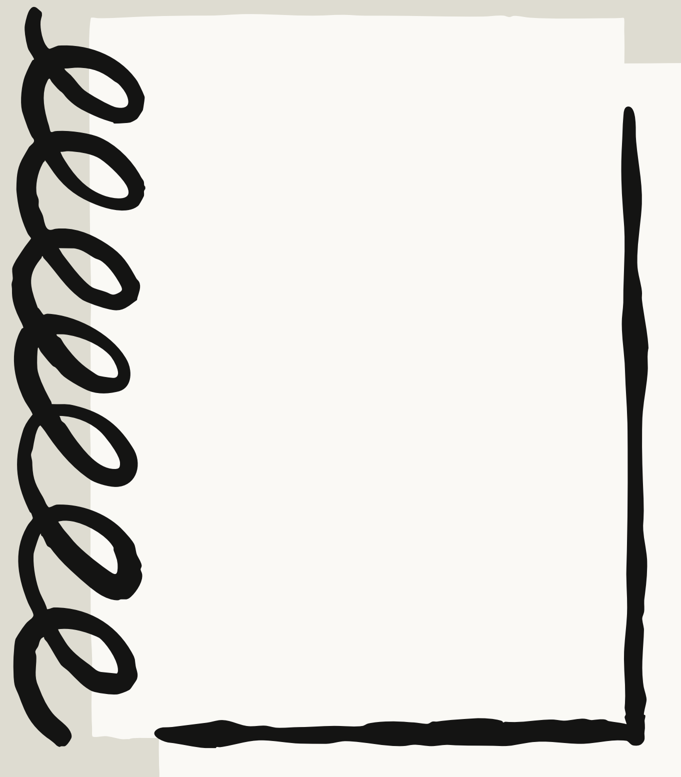

Chapter 9

Resources

Chapter 9

Resources

Customer stories

Cyera scales agentic AI across 1,500 employees with Claude Enterprise

Cyera on making Claude the front door to 40 tools (Q&A)

Cox Communications drives a 7x return on AI across its B2B funnel with Claude and Accenture

Workato connects enterprise AI agents to business systems with Claude

Full Anthropic customer story library

How Anthropic sells with Claude

How an Anthropic sales leader runs a 4,000-account book

Webinar: How Anthropic's sales team run their week with Claude

How Anthropic's business development team uses Claude to run inbound and outbound at scale

Webinar: Claude for Business Development Representatives

How one Anthropic seller rebuilt his team's workflows

How an Anthropic field marketer uses Claude Code to send weekly personalized updates to every sales rep

Set up and scale Claude for your organization

Anthropic's Sales plugin

Using Claude for sales account research

Scaling workflows with Claude at your organization

Claude Enterprise Administrator Guide

A guide to cost visibility and control in Claude

Deploying Claude Across Your Organization

Security and compliance documentation

Scaling workflows with Claude at your organization

Ready to roll out Claude to your revenue organization? Contact sales.

Claude

claude.ai