OutcomeCatalyst

A STARTER GUIDE

Claude for Healthcare Operations

The administrative side of the practice: denials, payer policy, contracts, credentialing. Nothing in here touches clinical care or patient treatment.

Written for
Practice managers, revenue cycle leads, COOs, CFOs

Time to read
About 20 minutes

Technical skill needed
None

outcomecatalyst.com
August 2026

WHAT IS IN HERE

Contents

Seven short sections. Read it straight through in twenty minutes, or skip to the part you need. Part 2 carries the rule about patient data, and it comes before any of the work.

1 **What Claude is, and what it is good at** 3
In plain terms, and what that means for the business side of a practice

2 **Getting set up** 5
Download it, pick a plan, and the one rule about patient data

3 **How to ask for things** 8
Four rules that change the answers, and three prompts you can copy

4 **Five jobs to try this week** 10
Coverage policy, a denial spike, payer rates, prior auth, credentialing

5 **Teaching it your way of working** 12
How to stop explaining yourself every time you open a new chat

6 **Where this stops** 14
What it will not do, and what to do about that

7 **What OutcomeCatalyst does** 15
How we take those four limits off the table, and when to call us

BEFORE YOU START

There is one app to download and no code anywhere in this guide. If you can install a program and write an email, you can do everything in here. The first job in Part 4 takes ten minutes and uses no patient data at all.

PART 1

What Claude is, and what it is good at

Claude is a chat window you can hand documents to. You type a question, you attach a file, and it answers. That is the whole idea.

The useful way to think about it here: **Claude reads the payer policy, the denial codes, and the contract, and tells you which of the three is actually causing the problem.**

That is the question nobody in a 40 person group has time to answer. One biller says it is a coding issue, another says the payer changed something in March, and the contract nobody has opened since it was signed says something different again. Working out which is true takes a week, so it does not get done, and the queue grows.

Give Claude ninety days of denial counts by payer and reason code, with no patient data in it, and ask what changed last month. Give it a coverage policy and ask which diagnosis codes it lists as supporting medical necessity. Give it a proposed fee schedule and your volume by code and ask what the new rates net you. Each of those is an afternoon of work and about ten minutes of reading.

It is not a database and it does not connect to your practice management system. It reads what you give it and answers what you ask. That is a smaller thing than the hype suggests and a much bigger thing than most people expect the first time they try it.

What it is genuinely good at on the business side

- **Reading policy and contract language and pulling out the specifics.** Coverage determinations, medical policy, participation agreements, fee schedule exhibits, prior authorization criteria.
- **Turning a mess into a table.** Three clearinghouse exports and two portal downloads, reshaped into the same columns so you can compare them.
- **First drafts of things nobody wants to write.** The appeal letter, the escalation to the provider representative, the board slide on where the money went.
- **Comparing things.** Proposed rates against current. This payer's policy against that one. What the contract says about a clean claim against what the payer is doing.
- **Asking better questions than you had time to think of.** Hand it the denial data and ask what it would look at next.

What it is not good at

- **Knowing anything about your practice.** It does not know your payers, your specialty mix, or your contracts until you tell it, in that conversation.
- **Being right about numbers without being checked.** It will produce a confident, well-written answer that is wrong, and it will look exactly like the ones that are right.
- **Codes from memory.** Code sets change every October and coverage policy changes more often. A code that came out of the model rather than out of a document needs verifying before it goes on a claim.

- **Doing anything in another system.** It cannot post an adjustment, work a claim, or update your practice management system. It writes, you act.
- **Anything clinical.** Nothing in this guide is a clinical tool. Coverage, contracts, and administration only. Care decisions stay with the people licensed to make them.

THE ONE HABIT THAT MATTERS

Check the numbers, and check the codes. Every time, at least by spot check. The rule people who use this well follow is simple: **it removes the typing, not the judgment.** A well-written appeal filed against the wrong policy version, or one day past the deadline, is money gone, and the fact that a machine typed it does not change who is responsible.

Where the time actually goes

The pattern in nearly every group is the same. A practice management system, a clearinghouse, four payer portals, a health record, and a folder of spreadsheets. None of them talk, so a person in the middle assembles the answer by hand, which is why it usually arrives after the month it was about.

THE THING THAT TAKES TOO LONG	WHAT CLAUDE CHANGES
Working out why denials spiked	A week of asking around becomes a ranked list by payer and reason code, with dollars attached
Finding what a coverage policy requires	Forty minutes on a government website becomes a plain answer with the codes and the documentation listed
Getting three payer exports into one shape	An afternoon of retyping becomes a first pass you correct
Modelling what a rate proposal nets you	A spreadsheet nobody has time to build becomes a bridge you can check line by line
Tracking credentialing expirations	A spreadsheet nobody updates becomes a dated worklist with lead times worked backwards

Two honest numbers, and one we will not give you

Banner Health, reporting on its own deployment, says **85 percent of surveyed staff saw significant time savings alongside improvements in work accuracy and workflow efficiency**. Read that carefully: it is a user survey self-report rather than an audited measurement, and Banner is a large health system rather than a 40 person practice. Directional, not a promise.

On the problem itself, HFMA's 2025 survey of **147 revenue cycle leaders** put commercial payer denials at **48.2 percent as the greatest threat to financial performance**, ahead of falling reimbursement at about 22 percent.

What we will not give you is a denial reduction percentage. Plenty circulate and none trace to anything we could stand behind. No software company has published an audited case study for a mid-market physician group, so anyone quoting a number for your practice is guessing.

PART 2

Getting set up

Ten minutes, start to finish. Get the desktop app rather than using the website, because the desktop app lets you drag files straight in.

STEP ONE, DOWNLOAD IT

claude.ai/download

Mac, Windows, iPhone, and Android. Free to download.

Download Claude

If that link does not open, type **claude.ai/download** into your browser. The same page is also at claude.com/download.

The five steps

1. **Download the app**
   Click Download under macOS or Windows, whichever you are on. You need macOS 11 or newer, or Windows 10 or newer.

2. **Open the file to install it**
   On a Mac it lands in your Applications folder. On Windows it lands in your Start menu.

3. **Sign in**
   Use your work email. If your organization already has Claude, use the account they gave you. That is the one compliance has looked at.

4. **Pick a plan**
   See the table below. Start on Pro, because every job in Part 4 that needs no patient data runs fine there.

5. **Drag a document in and ask it something**
   That is it. There is no setup wizard and nothing to configure.

PLAN	COST	WHO IT IS FOR
Free	Nothing	Fine for trying it on a public coverage policy. Limits will stop you partway through a long document.
Pro	$17 a month billed annually, or $20 monthly	**Start here.** Enough usage for real work, and it includes Projects, which matter in Part 5.
Max	From $100 a month	If you are running documents through it all day and Pro keeps stopping you.
Team	$20 a seat a month billed annually, minimum two seats	Shared settings and central billing, and conversations are not used to train the model by default. Still not eligible for a business associate agreement.
Enterprise	Published at $20 a seat plus usage at API rates. Check before you budget	The only plan with a HIPAA-ready option, plus audit logs and custom data retention. This is the conversation to have if patient data has to be involved.

The one rule about patient data

> THE RULE
>
> **Never paste patient identifiers into Claude unless your organization has a business associate agreement in place and has confirmed which product is covered by it.** Both halves matter. Having an agreement is not the same as the thing you are typing into being covered by it.

What an agreement can cover

Anthropic provides a business associate agreement covering its HIPAA-ready services, which means **the first-party API and Enterprise plans**. That is the whole list. Guides get this wrong constantly, so here it is plainly.

WHERE YOU MIGHT BE TYPING	CAN IT BE COVERED
Claude Enterprise	Yes, once accepted. Covers chat, projects, artifacts, file creation and code execution, voice, web search, research, and skills
Anthropic's first-party API	Yes, once signed and enabled
Free, Pro, Max, Team	**No.** Not eligible at all. Team's no-training default is a privacy improvement, not a HIPAA one
Cowork	**Explicitly excluded**, even on a covered account
The Office add-ins	**Explicitly excluded**, even on a covered account. That covers Claude for Excel, Word, PowerPoint, and Outlook
Outside connectors	Data flowing to a third-party service sits outside Anthropic's agreement. Separate conversation, separate vendor

Those last two catch people out, because Cowork and the Excel add-in are useful and sit right next to the covered surfaces in the same app. Being on Enterprise does not extend the agreement to them. If you are unsure which surface you are in, you are in the wrong one for patient data.

What the agreement does not do for you

- **It does not stop somebody pasting patient data into a personal account.** The biller with a free login at home is not covered by anything you signed. That is a policy and training problem.
- **It does not redact anything in an uploaded file.** If the PDF has a name and a date of birth in the header, the name and the date of birth went in. So did anything in a screenshot.
- **It does not produce your audit evidence.** Showing who accessed what, and proving it later, stays with you.
- **Access controls, workforce training, and minimum necessary stay yours.** Same as with any vendor you have ever signed with.

READ THIS BEFORE YOU TALK YOURSELF OUT OF THE WHOLE THING

The highest-value work in this guide uses no patient data at all. Coverage policy lookup. Denial code interpretation. Payer contract questions. Credentialing verification. Fee schedule work. Every one runs on policy documents, contracts, code sets, and de-identified counts. Start there and you have no exposure to manage and no agreement to wait on. Four of the five jobs in Part 4 sit in that category.

If you do go near patient data

De-identify before the model sees anything. Strip it at the source, in the export, before you attach the file. For denial work you rarely need more than payer, reason code, service line, count, and dollars, and that carries nothing about a person.

Never rely on the model to redact for you. Asking Claude to ignore the names in a PDF does not remove them, and asking it to redact a screenshot does not either. Redaction happens before upload or it does not happen.

One more, then go and see counsel. If your group handles substance use disorder records, a second federal rule stricter than HIPAA applies to them and **a standard business associate agreement does not automatically cover it**.

Two settings worth changing on day one

1. **Turn on file handling**
   Open **Customize** and switch on code execution and file creation. Part 5 needs it.

2. **Make a Project for revenue cycle**
   A folder with a memory. More in Part 5, but make one now.

If you cannot find a menu this guide mentions, it is almost always because file handling is off, or because you are signed in on the website rather than in the app. Check those two first.

PART 3

How to ask for things

Most people get mediocre answers because they ask mediocre questions. Four rules fix nearly all of it.

Rule 1. Say who you are and what the answer is for

Weak: "Why are we getting so many denials?"

Better: "I run revenue cycle for a four-site orthopedic group. Attached is ninety days of denial counts by payer and reason code, with no patient data in it. I meet our largest payer on Thursday and I need to know what to raise."

The second one works because it says what the answer is for. Same file, completely different output.

Rule 2. Give it the document, do not describe the document

Attach the actual policy PDF and the actual contract exhibit. Do not retype the sentence you think is the problem, because you have already made the judgment call about which sentence matters, and that judgment was the part you wanted help with. Attach the version with its date on it, too. Payer policy changes, and half the appeals that fail were written against the version that was live last year.

Rule 3. Say what the answer should look like

If you want a table, ask for a table and name the columns: payer, reason code, count, dollars, likely cause. If you want one page for a board packet, say one page. Unasked, you get prose, and prose is the hardest format to check.

Rule 4. Ask what it could not find, and what came from memory

Nobody does this and it is the most valuable of the four. End with: "List anything you could not find in the documents, anything you inferred rather than read, and any code or citation that came from your own knowledge rather than from the files I gave you."

That last clause is the healthcare-specific one. Codes and citations produced from memory are the most likely thing to be quietly out of date, and this is how you catch them before they reach a claim.

ONE MORE THING

Attach first, ask second. With a long document, putting your question at the end measurably improves the answer.

Three prompts you can copy

Use these as they are. Replace the italics with your own details.

PROMPT 1 · WORKING A DENIAL SPIKE

Attached is [ninety] days of denial data from our clearinghouse. It is de-identified: payer, denial reason code, service line, count, and dollars. There are no patient identifiers in it. You are a revenue cycle analyst. Build me a table of the largest denial buckets by dollars, with these columns: payer, reason code, what that code means in plain words, count, dollars, and the trend over the period. Then take the top three buckets and for each one give your best theory of cause, stated as one of three: our coding, a payer policy change, or a contract term. Say which evidence in the data points that way. Do not guess at anything outside the file. List separately anything you could not determine, and any code description you produced from your own knowledge rather than from the data.

PROMPT 2 · WHAT THE COVERAGE POLICY ACTUALLY REQUIRES

I need to know what Medicare's coverage policy requires for [service or test] in [state]. We keep getting denied for [medical necessity]. Tell me: whether there is a national coverage determination, which local coverage determination applies in that state, which diagnosis codes the policy lists as supporting medical necessity, what documentation it expects in the record, and anything that changed in the last thirty days. Cite the policy identifier and its effective date for every point. If you cannot find the policy, say so and stop rather than answering from memory.

PROMPT 3 · THE SECOND OPINION

You have just given me [the analysis above]. Now argue the other side. What is the strongest case that this appeal fails, or that I have misread which payer is the problem? What would the payer's own reviewer say? Be specific and use the documents and numbers I gave you. Do not hedge.

That third one is worth more than it looks. Left alone, Claude tends to agree with the framing you gave it. Asking it to argue against itself is how you find out your appeal rests on a paragraph that does not say what you thought.

PART 4

Five jobs to try this week

In order. The first takes ten minutes and needs no patient data, so start there.

1 Look up what Medicare's coverage policy in your state requires

10 MINUTES

What you need: the name of a service or test you keep getting denied on, and your state. **Zero patient data.**

What to do: open a new chat and use Prompt 2. If you added the CMS Coverage connector from Part 5, Claude reads the live coverage database rather than relying on memory.

What you get: what the policy requires, which diagnosis codes it lists as supporting medical necessity, what documentation it expects, and what changed recently.

Check before you act: open the policy itself and confirm the codes and the effective date. If Claude cannot name the policy and its date, treat what it said as a hypothesis and go look.

2 Work a denial spike

45 MINUTES

What you need: a denial export for the last ninety days, cut down to payer, reason code, service line, count, and dollars. **No patient identifiers.** No names, no account numbers, no dates of birth, no dates of service tied to a person. Strip those before you attach anything.

What to do: attach the export and use Prompt 1. Then take the largest bucket, attach the payer's policy and the relevant page of your contract, and ask which of the three is causing it. Run it on a month you have already worked by hand, so you can tell at once whether it got it right.

What you get: a ranked list of what changed, dollars on each line, a theory of cause for the top three, and the beginning of an appeal argument for the largest.

Check before you act: pull five real claims behind the biggest bucket and read them yourself. A pattern in the aggregate is a lead, not a finding. Then check the filing deadline before you write a single appeal, because the analysis is worthless if the window has closed.

3 Find out what a payer's proposed rates actually net you

45 MINUTES

What you need: the proposed fee schedule, your current one, and your volume by code. No patient data.

What to do: attach all three, ask for the net dollar effect by code and in total ranked by impact, then ask it to read the contract language and list the terms worth pushing back on.

What you get: a bridge from current revenue to what the proposal produces, and a list of terms to negotiate rather than a single yes or no.

Check before you act: recompute the top five lines yourself before quoting the total to anyone. Then read the language, not just the rates. Filing windows, recoupment periods, and how the contract defines a clean claim cost more than a few points on the conversion factor.

4 Build the prior authorization packet before you submit

20 MINUTES

What you need: the payer's policy for the service. If you want it checked against clinical documentation carrying patient identifiers, this needs an approved environment. Otherwise run it on the policy alone.

What to do: attach the policy and ask for a checklist of every requirement it states, in the order the payer looks for them, then ask what is commonly missing on submissions for this service.

What you get: a checklist for whoever assembles the packet, each requirement traced to the paragraph it came from.

Check before you act: this is an administrative completeness check and nothing more. It does not judge whether the service is appropriate for the patient, and must not be used that way. That decision stays with the ordering provider.

5 Get ahead of credentialing expirations

15 MINUTES

What you need: your provider roster with license, registration, board certification, and payer enrollment dates. Provider names are not patient data.

What to do: paste the roster, give today's date, and ask for everything expiring in the next 180 days as a dated worklist, with lead times worked backwards so you can see when each has to start.

What you get: a list by drop-dead date, with an owner column and a note on which lapses stop you billing.

Check before you act: verify every date against the board or the registry, not your own spreadsheet. Payer enrollment dates quietly cost the most.

IF YOU ONLY DO ONE

Do number one. Ten minutes, no patient data, nobody's approval needed, and you will know within the hour.

PART 5

Teaching it your way of working

By the third time you type out your payer list you will want to stop.

The easy way: a Project

A Project is a folder with a memory. Put things in it once, and every conversation inside that folder already knows them.

1. **Make a Project**
   Click Projects in the app, then create one. "Revenue cycle" and "Payer contracts" are where most people start. Free accounts get five, paid accounts get as many as you want.

2. **Put your standing context in it**
   Your payer list, your specialty mix, your top service lines, your appeal format. Whatever you explain twice.

3. **Work inside it**
   Every conversation there starts knowing all of it. You never type your payer list again.

WHAT TO PUT IN, AND WHAT TO LEAVE OUT

Put in only what Claude could not work out on its own. More material makes answers worse, not better, because the important part gets buried. And no patient data in a Project on a plan that is not covered, because a Project is a place things stay.

The stronger way: a skill

A skill is a saved instruction sheet Claude picks up on its own when it becomes relevant. Write down how you want a job done once, and from then on it does it that way without being asked. A Project is context inside one folder; a skill is a procedure that shows up wherever it is needed.

We wrote these for you

Five match the jobs in Part 4 and the work around them: **denials, payer contracts, prior authorization, credentialing, and clinic access.** Three more come with the pack: **practice operations, a compliance check, and quality measures.** They are free and yours to edit. **One file, one upload, all of them.** Claude opens whichever one your question needs and ignores the rest.

THE SKILLS PACK

outcomecatalyst.com/resources/healthcare-skills

One zip. Eight workflows. Free, no sign-up.

Download the pack

Inside: **oc-healthcare-toolkit.zip** to upload, each workflow on its own, and unzipped copies so you can read them before installing. They leave your payer list and thresholds blank, so the first run asks for them.

Installing the pack

1. **Open Customize**
   It is in the app sidebar. No Skills section there means file handling is off, so go back to Part 2.

2. **Click Skills, then the plus button**
   Then choose Create skill.

3. **Choose Upload a skill**
   Pick **oc-healthcare-toolkit.zip** from the pack.

4. **Check it is switched on**
   It appears in the list with a toggle already on.

5. **Go back to work**
   Attach a denial export, or ask about a coverage policy. Claude picks the workflow on its own.

Every skill says the same four things out loud

Every workflow carries the same four lines, so anything it produces says them on the face of the output: **this is administration support only**, not clinical advice. **Patient data belongs in an approved environment. Outputs get validated against your own policy and contracts**, not what the model remembers. And **the final decision goes to a qualified person**, named, every time. That is a feature, not boilerplate.

Letting Claude look things up for itself

Anthropic publishes hosted connectors that let Claude reach a data source directly, with no key and no setup beyond pasting in a web address. The most useful one here is **CMS Coverage**, for the Medicare Coverage Database: national and local coverage determinations, the codes each policy lists as supporting medical necessity, and what changed in the last thirty days.

To add it: open **Customize**, then **Connectors**, click the **plus** button, choose **Add custom connector**, and paste **https://hcls.mcp.claude.com/cms_coverage/mcp**. There are also connectors for the ICD-10 code set and the national provider registry.

> **BEING STRAIGHT ABOUT WHAT THESE ARE**
>
> Anthropic-hosted endpoints wrapping public government data. No key, no setup, no separate subscription. We will not call them free forever, because Anthropic publishes no rate limit and no service level commitment. Treat them as a good lookup, not as infrastructure.

This makes the safest work also the highest value: public policy lookup needs no patient data.

PART 6

Where this stops

Everything in this guide works. It also runs out, in the same place every time, and it is worth knowing where before you build a process on it.

You will notice it around the fourth or fifth week. The jobs in Part 4 work beautifully once. Then you want to run the denial analysis every month across every site, or you want two billers to do it and get the same answer, and something changes.

The four walls

- **It only sees what you paste.** A question that crosses the practice management system, the clearinghouse, and the payer portal means somebody pulls three exports and lines them up first. The gathering was the work.
- **The answer depends on who is typing.** Two billers work the same denial queue differently. One gets a good answer, one gets a different one, and nobody can say which is right or why.
- **Nothing gets written back.** Claude produces a worklist. Somebody still re-keys it into the practice management system, and the re-keying is where the errors come back in.
- **There is no trail.** No record of which policy version an appeal was written against, who ran the analysis, or when. Fine for a draft. Not fine when a payer disputes it eight months later, or when an auditor asks.

> READ THIS PART CAREFULLY
>
> None of that means chat is useless. The opposite. One person with the desktop app and twenty minutes of setup gets real, immediate value, and most practices never get that far. Start here. Get the value. Just know the shape of the wall before you run at it.

What comes after

Past those four walls, what you need is different in kind, not degree. Not a better chat window. A layer underneath that connects the systems you already run, so the answer arrives without anybody assembling it, comes out the same regardless of who asked, writes back where it belongs, and leaves a record.

That is what we build at OutcomeCatalyst. If you get to the point where the fourth wall is the one in your way, that is the conversation to have. There is no rush, and everything in this guide is worth doing on its own.

Start here

Take one service line you keep getting denied on and run Prompt 2 against it. Ten minutes, no patient data, and you will know.

PART 7

What OutcomeCatalyst does

The four walls in Part 6 are the reason we exist. Here is what we build, in the same plain terms as the rest of this guide.

We connect the systems you already run, so Claude works over your actual practice instead of over whatever somebody remembered to export.

Nothing gets replaced. Your practice management system stays. The clearinghouse stays. The health record stays where it is. We put a layer underneath that reads from all of them, resolves the same payer, provider, and service line across every one, and hands Claude a single connected picture. Then the questions that used to take a week take a sentence.

The four walls, and what removes each one

THE WALL	WHAT WE DO ABOUT IT
It only sees what you paste	We connect to your systems directly. The practice management system, the clearinghouse, the payer portals, the health record, the contract folder, and the spreadsheets people actually run the place on. If it has an interface, a database, or even just a monthly export, we can read it. Nobody assembles anything first.
The answer depends on who is typing	A workflow is defined once, with your payers and your rules in it, and it runs the same way for everyone. Two billers working the same denial queue get the same answer, because it is the same workflow rather than two people prompting.
Nothing gets written back	We write back. The worklist lands in the practice management system instead of being re-keyed into it. The flag reaches the person who can act on it while the filing window is still open. No re-keying, so no second round of errors.
There is no trail	Every output records what it read, which policy version, when, and who ran it. Role-based access, column masking, and a full audit trail come standard. When a payer disputes an appeal eight months later, the policy version it was written against is already recorded rather than reconstructed.

Reasoning over more than fits in a chat

A chat window can hold one month's denial export. It cannot hold three years across six sites and every payer you contract with. Ask a chat window which denials were never worked and aged past the filing deadline, and you will be uploading files until Thursday.

With the systems connected, that is one question against the whole book, and the answer arrives with the list attached. Same for which payer pays slower than the contract says, which service lines lose money after the last rate change, or which providers are approaching an enrollment lapse that will stop their claims. These are not harder questions. They are just questions nobody could ask before, because gathering the inputs cost more than the answer was worth.

THE PART PEOPLE ASK ABOUT MOST

The language model never does the arithmetic. Extraction, entity matching, joins, and every calculation run as governed code that produces the same result every time you run it. The model reads, explains, and drafts on top of that. So the numbers are repeatable and traceable to their source, and the writing is what the model is actually good at. Two different jobs, kept separate on purpose.

How it gets built

1. **We find the one that pays**
   A short look at how the work happens now, where the time goes, and which recurring decision costs the most. You get that answer whether or not you go further.

2. **We build one workflow**
   Not a platform, one named workflow: the data it starts from, the single thing it does, the artifact it produces, and the one person who owns it. Roughly three to four weeks with a clear before and after.

3. **It runs, and we keep it running**
   You keep the artifact and the value. We tune it and add the next one when you are ready.

A full install is typically live in about four to six weeks. A single workflow on top of an existing install is roughly three to four.

On security, plainly

Role-based access, column masking, and full audit trails come standard. For the most sensitive environments there is a locally hosted option that keeps everything on your own infrastructure, where nothing leaves your walls. We are HIPAA-aligned and SOC 2 Type 2 aligned, with the formal audit underway and expected to complete before year-end.

When to call us

Not yet, if you have not done Part 4. Go run the coverage policy lookup on a service you keep getting denied on. Everything in this guide is worth doing on its own and costs you nothing but an afternoon.

Call us when you have hit a wall in Part 6 and can name it. That is the conversation worth having, and by then you will know exactly which one it is.

START WITH THE WEDGE

outcomecatalyst.com

Point us at the process that costs you the most time or money, and we will come back with the workflow, the owner, and the ROI.

Book a call

THE SMALL PRINT

Credits and disclaimers

WHAT WE USED	WHAT WE TOOK FROM IT
Anthropic claude.com, support.claude.com, privacy.claude.com	Plans, prices, download instructions, system requirements, settings paths, connector addresses, data handling, and the business associate agreement coverage and exclusions in Part 2. Checked against live pages in August 2026.
OutcomeCatalyst outcomecatalyst.com/resources/healthcare-skills	The eight healthcare operations skills in Part 5. Released under the Apache 2.0 license, which is included with the files. Free to use, change, and pass on.
Banner Health	The 85 percent figure, which is Banner's own published survey of its staff, reproduced with its original qualifier.
HFMA	The 48.2 percent and roughly 22 percent figures, from a 2025 survey of 147 revenue cycle leaders.

This document is general information about administrative and revenue cycle work. It is not legal, compliance, coding, billing, tax, or clinical advice, and it does not create a professional relationship. Nothing in it is intended to inform diagnosis, treatment, or any decision about the care of a patient. Coverage rules, medical policy, coding guidance, contract terms, and privacy obligations vary by payer, by state, and by the wording of your own agreements, and they change frequently. What you may put into any tool is governed by your organization's agreements and policies rather than by this guide, so read the agreement in front of you and ask counsel or your privacy officer before patient data goes near a chat window. Records subject to the federal substance use disorder confidentiality rule carry a stricter standard than HIPAA, and a standard business associate agreement does not automatically cover them. Every workflow here produces a draft for a qualified person to review. None produces a decision, a code you may bill without verification, or a clinical judgment. Product details about Claude, including which plans and surfaces a business associate agreement can cover, were verified against Anthropic's pages in August 2026 and change frequently, so confirm them for your own organization before relying on them. The Banner Health and HFMA figures are self-reported and survey-based rather than audited, and are reproduced with their original qualifiers. No software company has published an audited case study for a mid-market physician group, and we publish no denial reduction figure because none we found is supportable. The skills we provide are released under the Apache 2.0 license and come with no warranty of accuracy or fitness for your purpose. Skills can run code on your machine with your privileges, so read anything you install, including ours. Unzipped copies are included for that reason. Claude is a product of Anthropic. This is an independent guide, not affiliated with or endorsed by Anthropic, any payer, or any practice management, clearinghouse, or health record vendor named or implied. All trademarks belong to their respective owners.