THE CHAIRSIDE AI KIT

Day-1 Install & Run Guide

A cookbook for the eight skills.
Zero coding. 30 minutes to first return.

SESSION C COMPANION

BULLETPROOF DENTAL CE SUMMIT · AUG 9, 2026

BUILT BY AIPG · BLAKE MCCLELLAN

Front matter

How to use this guide

You are holding a 40-page manual. Do not read it front to back.

Do this instead:

1. **Read Chapters 1–3 once.** Fifteen minutes. That's the concept, the prereqs, and the universal install.
2. **Pick one skill.** Chapter 4 (Money-Leak Finder) if you're following Blake's Monday recommendation.
3. **Read that skill's chapter, install it, run it.** Forty minutes, start to finish.
4. **Come back later** for the other seven — one a week is a good pace.

The 90-day rollout plan (Chapter 12) tells you what order and pace works for most practices. Follow it or don't.

> **THE PROMISE**
>
> 15 minutes for one skill. 45 minutes for all eight. 5 minutes to first return (Money-Leak Finder, Monday morning).
>
> Not a set-it-and-forget-it tool. Every output is a starting point the doctor reviews before it goes to a patient, a staff member, or a vendor. **AI drafts. Humans decide.**

Contents

Table of contents

Pages are approximate. Each skill chapter is roughly 3 pages.

Chapter 01

What the Chairside AI Kit is

Eight pre-built AI skills for dental practices, packaged for one-click install into Claude Desktop, with paste-in versions for ChatGPT. Free. Zero lock-in.

The one-sentence version

The Chairside AI Kit is 8 pre-built AI skills for dental practices, packaged for one-click install into Claude Desktop, and available as paste-in versions for ChatGPT. It's free, it's yours to keep, and there is nothing to buy or subscribe to.

What's inside

#	SKILL	SOLVES	TIME TO RUN
1	**Money-Leak Finder**	Where is money bleeding out of subscriptions and vendor charges?	5 min
2	**Case Closer**	Turn a treatment plan into a patient-ready handout	15 min
3	**Fee Auditor**	Which CDT fees are underpriced, and what should I raise them to?	20 min
4	**EOB Decoder**	What does this EOB actually say, and what does the patient owe?	5 min / EOB
5	**Cyber Self-Audit**	Am I ransomware-ready? What's my HIPAA exposure?	15 min
6	**Morning Huddle**	Give me tomorrow's huddle sheet from today's schedule	3 min
7	**Hiring Kit**	Job post + interview scorecard + 30/60/90 onboarding for one role	25 min
8	**Vision Coach**	I have to have a hard team conversation. Walk me through it.	20 min

What a "skill" actually is

If you didn't sit through Session A: a **skill** is a pre-loaded set of instructions that tells Claude how to handle a specific task the same way every time.

- **Without a skill:** you type a request, Claude does its best with whatever context it can guess.

- **With a skill:** Claude follows a structured script — same steps, same guardrails, same output format — every time you run it.

Think of it as a *recipe* loaded into Claude. You give it the ingredients (your data), it gives you the finished dish (the report, the handout, the huddle sheet). You install a skill once. You run it dozens of times.

How the 8 skills fit together

The kit maps to the eight most repetitive, non-clinical parts of running a dental practice.

MONEY-FACING

- Money-Leak Finder → find hidden savings
- Fee Auditor → capture what you're underpricing
- Case Closer → turn plans into decisions
- EOB Decoder → resolve insurance faster

OPS-FACING

- Morning Huddle → run the day
- Cyber Self-Audit → protect the practice
- Hiring Kit → hire the right person
- Vision Coach → lead through the hard moments

Together, they cover roughly 60–70% of the ops surface of a solo or small-group dental practice. Not clinical work. Not the front-desk phone script. Not treatment planning itself. But everything around those — the plumbing that eats your evenings.

What this kit is NOT

- **Not a replacement for your practice management software.** Dentrix, Eaglesoft, Curve — those systems are the source of truth. The skills read data *from* them; they don't replace them.
- **Not for clinical decisions.** No diagnosis. No treatment planning. No prescription guidance. That's what dental school and your CE hours are for.
- **Not HIPAA-compliant for patient-identifying workflows.** These skills run on consumer Claude and consumer ChatGPT. That means you must strip patient identifiers *before* pasting

anything. Every patient-adjacent skill has built-in guardrails that warn you if you slip an identifier in — but the responsibility is yours.

- **Not a "set it and forget it" tool.** Every output is a *starting point* the doctor reviews before it goes to a patient, a staff member, or a vendor. AI drafts; humans decide.

Chapter 02

Prerequisites

What you need before Monday morning. Fifteen minutes of setup, one time.

Software

- **Claude Desktop** (free tier works fine) — download at `claude.ai/download`. Mac (Apple Silicon and Intel), Windows (standard installer, no admin rights needed). Alternative: `claude.ai` in your browser — same functionality.
- **A modern web browser** for viewing the reports Claude renders — Chrome, Safari, Firefox, or Edge.
- **A CSV editor** — Excel, Numbers, or Google Sheets. You already have one.

Optional — but strongly recommended

- **Claude Pro** ($20/mo) if you plan to run the kit more than 2–3x/week. Free tier has message limits that get hit fast if you're using it for real practice work.
- **A dedicated computer** for practice AI work — even an old laptop. Reason: you turn off "training on my data" once, and that setting only applies to that specific account on that specific device. A dedicated login = a dedicated posture.

Where to download the kit

> **PLACEHOLDER — RESOURCE URL TBD**
>
> Blake will confirm the resource page URL by summit day. Likely candidates: `aipg.co/kit` or `bulletproof2026.com/kit`. The QR code on Slide 9 points to the same location.

You'll get a single ZIP file. Unzip it. You'll see:

```
chairside-ai-kit/
  chairside-money-leak-finder/
  chairside-case-closer/
  chairside-fee-auditor/
  chairside-eob-decoder/
```

```
chairside-cyber-self-audit/
chairside-morning-huddle/
chairside-hiring-kit/
chairside-vision-coach/
gpt-fallback/
README.md
install-guide.pdf  (this document)
```

The ChatGPT fallback bundle

If your practice is standardized on ChatGPT instead of Claude — no problem. The `gpt-fallback/` folder contains a paste-in version of each of the 8 skills, formatted as single prompts you paste into a new ChatGPT chat.

How the ChatGPT versions differ from the Claude versions:

- They're collapsed into one long prompt per skill (Claude skills are multi-file).
- File-handling is limited — you'll paste data as text more often (that's fine; every skill has a paste-as-text path).
- There's no code sandbox for math — the ChatGPT version does the arithmetic inline, which is slightly less reliable but still workable.
- **Verdict:** Claude is the better home for these, but the ChatGPT versions work on free ChatGPT with no plugins.

Grok, Gemini, and other model users: the ChatGPT fallback prompts work in almost any modern LLM chat interface. Paste and go.

The privacy setup you do ONCE

Before you run any skill, do this (5 minutes, one time):

> **CLAUDE DESKTOP**
>
> 1. Open Claude Desktop
> 2. Click your name/profile in the bottom-left
> 3. Settings → Privacy
> 4. Turn OFF "Improve Claude for everyone" (or however it's currently labeled)
> 5. Confirm you're using YOUR own logged-in account, not a shared one

CHATGPT

1. Open ChatGPT
2. Click your profile in the bottom-left
3. Settings → Data Controls
4. Turn OFF "Improve the model for everyone" (or "Chat History & Training" depending on account type)

Do this on every device you'll use for practice work. That's it. Now the data you paste never gets fed back into future training runs.

One more privacy note

Even with training off, treat every paste as if a stranger might see it. That's why every patient-adjacent skill has a **de-identification gate** built in. Strip names, DOBs, member IDs, and full addresses *before* pasting. The skills catch what they can, but the human is the last line of defense.

Chapter 03

Universal install — how to install any skill

You only learn this once. Every skill installs the same way.

Step-by-step: install a skill

Step 1 — Unzip the kit

Download `chairside-ai-kit.zip`. Double-click to unzip. You'll get a folder called `chairside-ai-kit/` on your desktop or Downloads folder.

> SCREENSHOT *A Finder window (Mac) or File Explorer window (Windows) showing the unzipped folder with 8 sub-folders visible — one per skill.*

Step 2 — Open Claude Desktop

Launch Claude Desktop. If you haven't downloaded it yet: `claude.ai/download` → install → sign in with the account you'll use for practice work.

> SCREENSHOT *The Claude Desktop launch screen — a chat window with the input bar at the bottom.*

Step 3 — Open Skills settings

Click your name/profile in the lower left → **Settings** → **Skills** tab (if you don't see a Skills tab, make sure you're on the latest version of Claude Desktop — it may be under "Advanced" on older builds).

> SCREENSHOT *Settings modal open, "Skills" tab highlighted in the left sidebar. Main pane shows an "Import Skill" button and any already-installed skills.*

Step 4 — Import the skill

Click **Import Skill** (or the "+" button, depending on your version). A file-picker opens. Navigate to the unzipped kit folder → pick the sub-folder for the skill you want (e.g., `chairside-money-leak-finder`) → confirm.

Alternatively, some Claude Desktop versions accept a `.zip` of the skill folder — if the folder-import doesn't work, right-click the skill's folder → "Compress" → import the `.zip`.

> **SCREENSHOT** *File picker dialog with `chairside-money-leak-finder/` folder selected, "Open" button highlighted.*

Step 5 — Save and verify

The skill appears in your Skills list with its name (e.g., "Chairside Money-Leak Finder") and a short description. Close Settings.

> **SCREENSHOT** *Skills tab showing "Chairside Money-Leak Finder" now in the list, with a description underneath. A toggle switch shows it as "Enabled."*

Step 6 — Run it once

In a fresh Claude chat, type: **"Run Chairside Money-Leak Finder"** (or just "Money-Leak Finder" — Claude will figure it out).

The skill loads. It'll speak first — usually asking about privacy setup and what data you'll provide. Follow its lead.

> **SCREENSHOT** *Claude Desktop chat window with a user message "Run Chairside Money-Leak Finder" and Claude's response beginning: "Before we start — a note on privacy..."*

Common gotchas

> **"I DON'T SEE A SKILLS TAB."**
> Your Claude Desktop is out of date. Go to `claude.ai/download` and reinstall. If Skills is still missing, use the browser version at `claude.ai` — the browser version always has the latest features.

> **"IMPORT FAILED."**
> Two common causes:
>
> 1. You tried to import a `.md` file instead of the folder — import the whole folder.
> 2. The folder is nested one level too deep. Open the folder — if you see another folder with the same name inside, import THAT inner folder.

> **"THE SKILL INSTALLED BUT DOESN'T RUN WHEN I SAY THE NAME."**
>
> Make sure it's enabled (toggle in Skills settings). Type the exact skill name from the Skills list. Or say "Run the money-leak finder skill" — Claude interprets loosely.

> **"I USE CLAUDE IN THE BROWSER, NOT DESKTOP. DOES THIS WORK?"**
>
> Yes. The steps are identical — Skills settings are in the same place (profile → Settings → Skills). Import works the same way.

> **"I DON'T HAVE A SKILLS TAB EVEN IN THE BROWSER."**
>
> Skills is a paid-tier feature on some plans. Check `claude.ai/pricing`. Free tier may not support Skills yet — in which case, use the ChatGPT fallback bundle. Same skills, just paste-in.

How to verify the install worked

1. In a fresh chat, type: **"What skills do you have available?"**
2. Claude lists installed skills. You should see the one you just added.
3. Type: **"Run [skill name]."**
4. The skill's opening message should appear — for most skills, it starts with a privacy/de-identification note.

If both of those succeed, you're good.

Installing all 8 at once

Repeat steps 3–5 for each skill folder. Import all 8 in one sitting — takes about 5 minutes total. You do NOT have to install all 8 to get value from any one. Install one, run it, and come back for the next when you're ready.

Money-Leak Finder

Audit a bank/credit-card CSV and get a prioritized list of recurring charges that look like waste — with confidence scores and ready-to-send renegotiation scripts.

WHAT IT DOES

Money-Leak Finder audits your practice's bank or credit-card statement (as a CSV) and returns a prioritized list of recurring charges that look like waste — duplicate subscriptions, overpriced print/internet, zombie SaaS, quarterly overages. Along with each finding: a confidence score, the reason it flagged it, and a ready-to-send renegotiation or cancellation script.

What most practices find: $500–$1,500/month in leaks they didn't know existed. That's $6,000–$18,000/year uncovered in a 5-minute run.

BEST-FIT SCENARIO

Reach for this skill when:

- You have that nagging feeling your subscriptions are out of control but you can't remember what you signed up for
- You just onboarded a new office manager and want a clean baseline
- Your bookkeeper flagged "miscellaneous charges" or you've had a print-lease invoice creep on you
- You're heading into budget season and need to cut fat without cutting people

Frequency: run this quarterly. Once every 3 months is enough to catch new leaks; more often is diminishing returns.

INPUTS (WHAT YOU PROVIDE)

- **A CSV of 60–90 days of bank or credit-card transactions.** Export from your bank or card portal — Transactions/Activity → Export → CSV.
- **Critical:** delete the full account number / card number column before uploading. The skill never needs it.
- **What the CSV should contain:** date, vendor/description, amount. That's it. If your bank exports extra columns (category, balance), the skill parses generously.

- **What it should NOT contain:** patient names, patient charge details, or anything that would identify a specific patient.

Why 60–90 days: catches quarterly charges (print overages, annual license renewals) that a 30-day export would miss.

OUTPUTS (WHAT CLAUDE PRODUCES)

- **Top-line finding:** "Found *$X/mo (~$Y/yr)* in likely leaks across N charges."
- **Ranked table:** Vendor · Category · $/mo · Action (Cancel / Keep / Renegotiate) · Confidence · Why.
- **"This week" shortlist:** the 3 highest-dollar moves, each with a pre-filled script attached.
- **Periodic charges to review:** e.g., a quarterly print overage.
- **Context:** total leak vs. the $600–800/mo benchmark and as a % of overhead.

The output renders as a scannable HTML report (Artifact) — you can print it, save it as PDF, or share it with your bookkeeper.

THE EXACT PROMPT TO TRIGGER IT

> **PROMPT**
>
> *Run Chairside Money-Leak Finder. I'll upload a 60-day bank statement CSV.*

Then follow its lead — it will confirm privacy, ask for the CSV, and run.

WORKED EXAMPLE

> **WORKED EXAMPLE**
>
> **Practice:** "Bright Smile Dental," a 2-op solo practice in Oklahoma City. Annual collections ~$1.2M. Owner: Dr. Marcus Cole.
>
> **Input:** 87 days of Chase business-checking CSV. 142 transactions. Card number column deleted before upload.
>
> **OUTPUT THE SKILL RETURNS (EXCERPT)**
>
> **Found $1,180/mo (~$14,160/yr) in likely leaks across 9 charges.**

VENDOR	CATEGORY	$/MO	ACTION	CONF.	WHY
Verizon Business Fiber	Internet	$340	Renegotiate	High	Local benchmark for 500Mbps in OKC is ~$180. Paying nearly 2x.

VENDOR	CATEGORY	$/MO	ACTION	CONF.	WHY
Konica Minolta Lease	Print lease	$285	Renegotiate	High	Quarterly overage of $190 — printing 2x contracted volume.
Dropbox Business (14 seats)	SaaS	$210	Verify	High	Team is 6 people. 14 seats is 8 more than headcount.
Shred-It	Compliance	$180	Renegotiate	Medium	Market rate for once-monthly shredding at practice size is ~$85.

THIS WEEK — TOP 3 MOVES

1. Call Verizon Business. *Script:* "I'm reviewing our business services and I see we're paying $340/mo for 500Mbps fiber. Comparable local providers are quoting $180 for the same speed. I need this brought down to at least $200 or I'll need to switch. What can you do?"

2. Audit Dropbox seat count. *Script:* Log into Dropbox admin → Users → downgrade to actual headcount (~$70/mo saved).

3. Call Konica. *Script:* "Our print overage last quarter was $570. Our monthly contract is $285. I need our monthly volume adjusted up to match actual usage, or I need to explore a lease buyout. What are the options?"

What Dr. Cole does next: Makes the three calls between patient blocks the next day. Saves $605/mo by end of week ($7,260/yr). Deletes the extra Dropbox seats himself.

TROUBLESHOOTING

"THE SKILL SAYS IT CAN'T DETECT COLUMNS."

Your bank's CSV format is quirky. Open the CSV, look at the first row (headers), and tell the skill: "Column A is date, column B is description, column C is amount." Rerun.

"IT FLAGGED SOMETHING AS WASTE THAT ISN'T."

Confidence scores matter — "Verify" means the skill isn't sure. Don't cancel it until you've confirmed. Every finding has a "Why" column — read it. If the reasoning is wrong, correct the skill: "That's actually X and Y, not a duplicate. Rescore."

"THE RENEGOTIATION SCRIPT DOESN'T FIT MY SITUATION."

Say: "Rewrite the script assuming I've been a customer for 8 years and I'm not willing to switch — I want a loyalty discount, not a switch threat."

"MY STATEMENT INCLUDES PAYMENT PROCESSOR FEES."

Include them. The skill knows those aren't leaks — they're revenue-attached costs, and it'll skip them.

"WHAT ABOUT PAYROLL OR RENT?"

Skip those. Money-Leak Finder is for discretionary and recurring vendor charges. Payroll, rent, insurance, taxes go in a different bucket (accountant's job).

SUCCESS SIGNAL

SUCCESS SIGNAL

You know it's working when:

- You've cut at least $300/mo of confirmed waste within 2 weeks of running it
- You've had at least one "wait — we're paying for THAT?" moment
- You're setting a calendar reminder to run it again in 3 months

If you found nothing, either your practice is already lean (rare) or your CSV export was incomplete. Re-export with more history and try again.

Case Closer

Take a de-identified treatment plan and return a patient-ready case presentation with tiered financing, phasing, and a 4-touch follow-up sequence. Anchored to the patient's stated goal, never upsold past it.

WHAT IT DOES

Case Closer takes a **de-identified** treatment plan and turns it into a patient-ready case presentation — plain-language "why," tiered financial menu with payment options, illustrative monthly estimates, phasing, and a Day 0 / Day 3 / Day 7 / Day 14 follow-up sequence. Anchored to the patient's stated goal, not upsold past it. The doctor reviews before any patient sees it. The skill drafts; you decide.

BEST-FIT SCENARIO

- A patient just accepted a complex plan chairside, but you know they need to see the money and phasing in writing before they commit
- Your team is asking for a case-presentation template that doesn't feel salesy
- You want to standardize how tx plans get presented — same voice, same math, same disclosures every time
- You're onboarding a new TC and need a scaffolding

Frequency: every time a plan is above ~$3K and the patient needs a takeaway. For solo practices that's usually 2–5x/week.

INPUTS (WHAT YOU PROVIDE)

Strip these BEFORE pasting (the skill halts if it detects an identifier):

- Patient name, DOB, address, phone, email
- Member ID, subscriber ID, claim number
- Insurer's identifying group name (a generic "PPO Basic 2024" is fine)

Keep and share:

- Procedures + fees (tooth numbers are fine)

- Each item flagged *urgent* (pain / infection / can't chew) or *recommended* (needed but can wait)
- One-line clinical "why" per treatment area, plain language
- Optional: age band ("late 50s"), insurance type, remaining annual max, rough coverage %, patient's goal in their own words
- Your practice's financing partners (name, term, promo vs. APR)

One-time brand setup (do this once, reuse forever): Create a Claude Project called `[Your Practice] — Patient Materials`. Add a brand config file with practice name, hex colors, Google Font names, phone, booking URL, financing partners, and voice guidelines. Add your logo. Every case presentation you generate inside this project inherits the brand — zero re-setup.

OUTPUTS (WHAT CLAUDE PRODUCES)

- **Three-tier financial menu:** Comprehensive / Phased / Essential-now. Each with subtotal, insurance offset, patient portion, illustrative monthly estimate range labeled "estimate, subject to approval."
- **Phasing breakdown:** which procedures happen when
- **Compliant financing language:** no promises of approval, no credit-score references, deferred-interest disclosure on any promo
- **Banned words check:** no "best, guaranteed, painless, permanent" anywhere
- **The handout Artifact:** a printable one-page (or two-page) case presentation, on-brand
- **The 4-touch follow-up sequence:** Day 0, 3, 7, 14 — de-identified, `[First name]` placeholders
- **Internal coaching note for the team:** which tier to lead with, who owns the follow-up, the cost-of-waiting line

THE EXACT PROMPT TO TRIGGER IT

> **PROMPT**
>
> *Run Chairside Case Closer. I have a de-identified treatment plan to present. Late-50s patient, in-network Delta PPO, remaining annual max $850, coverage 80/50. Their goal in their words: "I want to be able to chew steak with my wife again." Here are the procedures.*

Then paste the plan with tooth numbers, procedures, fees, and urgent-vs-recommended flags.

WORKED EXAMPLE

WORKED EXAMPLE

Practice: "Sunrise Family Dental." **Patient goal:** "I want to chew again — the pain is why I stopped eating with my family."

INPUT (PASTE TO CLAUDE)

De-identified plan:

- #30 — root canal ($1,450) — URGENT (currently in acute pain)
- #30 — porcelain crown ($1,850) — URGENT (paired with RCT)
- #14, #15 — two composite fillings ($480) — recommended (early decay)
- #19 — implant + abutment + crown ($5,600) — recommended (missing 6 mo, chief concern)
- Full-mouth periodontal maintenance every 3 mo — recommended

Insurance: Delta Dental PPO, remaining annual max $850, 80% basic / 50% major.
Patient goal (their words): "I want to chew steak with my wife again." Age band: late 50s.
Financing: CareCredit (24 mo no-interest promo, longer-term APR), Sunbit.

WHAT CLAUDE PRODUCES (EXCERPTS)

Top of handout: "What we heard from you: 'I want to chew steak with my wife again.'"

Option A — Get you eating first (Essential now): $3,780 total. Insurance estimate: –$680. Patient portion: **$3,100**. Illustrative monthly with CareCredit 24-mo promo: ~$130/mo (estimate, subject to approval). Covers: RCT and crown on #30, plus the implant on #19 — the two things standing between you and steak.

Option B — Phased care over 6 months: $9,380 total. Insurance estimate: –$850. Patient portion: **$8,530**. Illustrative monthly with 60-mo APR: ~$180/mo. Covers everything in A plus the two composites and a year of maintenance.

All monthly figures are estimates, subject to lender approval. Insurance figures are estimates, not guarantees.

THE FOLLOW-UP SEQUENCE

Day 0 (text): "[First name], really appreciated you coming in today. Here's the handout Dr. [Name] walked through — [link]. Any question you didn't want to ask in the chair, text us here."

Day 3 (email): "[First name] — quick check-in. We know steak is the goal. If Option A feels like the right first step, we can get #30 handled next Thursday. Want us to hold that slot?"

(Day 7 and Day 14 continue.)

INTERNAL COACHING NOTE FOR THE TEAM

Lead with Option A. Patient's stated goal is functional chewing — a comprehensive pitch will feel like an upsell. If they ask about #19, that's an enthusiastic yes. If they ask about composites, mention they can wait until the plan resets. Owner of follow-up: Sarah, TC. Cost-of-waiting line: "Every month the RCT waits, the tooth risks non-restorability."

TROUBLESHOOTING

"IT STOPPED AND SAID IT DETECTED AN IDENTIFIER."

The de-ID gate caught something — a name, DOB, or ID. Look at your paste. Scrub. Re-paste. Do NOT ask Claude to "ignore it and continue" — the halt is intentional.

"THE MONTHLY ESTIMATE IS WILDLY OFF."

Give the skill more info: "CareCredit 24-mo promo has a $2K minimum for promo pricing. Below that, it's APR-only." Rerun. The math updates.

"THE HANDOUT DOESN'T LOOK ON-BRAND."

Set up the Project + brand config (see Inputs). Case Closer respects it. Without the Project, you'll re-set brand every session.

"THE THREE-TIER MENU FEELS TOO COMPLEX."

Ask: "Simplify to two tiers — Essential now and Comprehensive. Skip Phased." Case Closer adapts.

"IT REFUSED TO WRITE IN THE TONE I ASKED FOR."

It won't use "best" / "guaranteed" / "painless" / "permanent" — banned words. Rephrase. It will use anything else.

SUCCESS SIGNAL

SUCCESS SIGNAL

- Your team stops writing case-presentation handouts from scratch
- Your case acceptance rate for plans >$3K measurably improves in 60–90 days
- Patients start referencing specific parts of the handout when they come back to book

> If patients are ignoring the handout entirely, the goal-anchoring isn't working — capture better patient-goal language at the exam.

Chapter 06 · Skill 3 of 8

Fee Auditor

A 5% fee increase is roughly a 14% profit increase. Fee Auditor tells you which CDT codes are underpriced and by how much — with talking points for your team and a re-key CSV for your PMS.

WHAT IT DOES

Fee Auditor takes your fee schedule (CDT code · description · current fee), your ZIP code, and your overhead %, and returns a profit-leverage analysis + a list of flagged underpriced codes + a tiered recommended increase + a re-key CSV you paste back into your PMS.

The core insight: at 65% overhead, a 5% fee increase is roughly a **14% profit increase**. Your fixed costs don't move when fees do — every new dollar of collected fee is nearly pure margin.

BEST-FIT SCENARIO

- You haven't raised fees in 12+ months (most practices are behind by CPI + market rate)
- You're heading into January and want to time an increase to the new plan year
- You just took over a practice and don't know if the previous owner's fees are anywhere near market
- You're thinking about going out-of-network on a PPO and need to know what your "real" fees should be first

Frequency: annually. Same month every year. Most practices do it in October–November to prep January increases.

INPUTS (WHAT YOU PROVIDE)

- **Fee schedule** — CSV or XLSX export from your PMS. Format: `CDT code · description · current fee`. Or paste the codes directly.
- **Your ZIP code** — helps calibrate cost-tier context (90210 benchmarks differently than 74133).
- **Overhead %** — total practice overhead as a percent of collections. Default: 65%.
- *Optional:* annual collections ($1.8M etc.) — unlocks a real-dollar profit projection.
- *Optional:* year fees were last updated (2022 etc.) — unlocks the inflation-gap calculation.
- *Optional:* in-network yes/no — unlocks a PPO note (realized vs. billed).

OUTPUTS (WHAT CLAUDE PRODUCES)

- **Headline:** "A [X]% fee increase is a [Y]% profit increase for your practice"
- **Profit-leverage snapshot:** overhead, +5% and +10% profit lift in % and $, stale-fee gap since last update
- **Flagged codes table:** likely underpriced codes with directional range, gap %, confidence, basis — with "Directional only — confirm vs. Patterson/FAIR Health" stamped on every benchmark
- **Recommended increase schedule:** CDT · Description · Current · Recommended · % Increase · Tier (baseline 4.5% all codes; targeted 8–12.5% for flagged codes)
- **Plain-English math** so you can explain it to your accountant or your spouse
- **Talking points:** team briefing, front-desk patient response ("Did prices go up?"), annual-cadence argument
- **Re-key CSV:** downloadable, formatted for PMS import

THE EXACT PROMPT TO TRIGGER IT

> **PROMPT**
>
> *Run Chairside Fee Auditor. I'll upload my fee schedule as a CSV. My ZIP is 74133 (Tulsa suburb), overhead is 62%, collections are $1.6M last year, we last updated fees in 2023, we're in-network with Delta and MetLife.*

WORKED EXAMPLE

> **WORKED EXAMPLE**
>
> **Practice:** "Sunrise Dental Group." 25 CDT codes. Last update: 2022. Collections: $1.8M. Overhead: 65%.
>
> **WHAT CLAUDE PRODUCES (EXCERPTS)**
>
> **A 10% fee increase is a +29% profit increase for your practice.**
>
> - Your overhead: 65%
> - +5% fees → +14% profit / +$90,000 top-line
> - +10% fees → +29% profit / +$180,000 top-line
> - Stale-fee gap: 12.4% behind CPI since 2022

FLAGGED CODES (LIKELY UNDERPRICED — DIRECTIONAL)

CDT	DESCRIPTION	CURRENT	DIRECTIONAL RANGE	GAP%	CONF.
D2740	Porcelain crown	$985	$1,150–$1,350	–16.8%	High
D4341	Perio scaling 4+ teeth	$195	$230–$275	–18.3%	High
D0210	Full-mouth series	$110	$140–$170	–27.3%	Med

(Directional only — confirm vs Patterson Dental Fee Survey or FAIR Health before adjusting insurance contracts.)

RECOMMENDED INCREASE SCHEDULE

- **Tier 1 (baseline, 4.5% on all codes):** covers CPI catch-up
- **Tier 2 (targeted, 8–12.5% on 6 flagged codes):** captures market gap on high-leverage codes

TALKING POINT — "DID PRICES GO UP?"

"Yes, this January we adjusted our fees to reflect the increased costs of materials, wages, and overhead. It's the first time we've adjusted since 2022. The good news is we're still one of the more affordable practices in the area — most others adjust every year."

TROUBLESHOOTING

"IT ONLY PARSED 12 OF MY 60 CODES."

Your CSV headers may be non-standard. Open it, name the columns `CDT`, `Description`, `Fee`, and re-upload.

"THE BENCHMARKS FEEL TOO LOW."

They're directional floors — the skill errs conservative. If Patterson data says higher, use that. The skill's benchmarks are a "you're at least this low" signal, not a "raise to exactly this" recommendation.

"I'M IN-NETWORK. HOW MUCH SHOWS UP?"

Depends on your PPO contract mix. Realized increase on PPO patients is what the contract's allowed amount permits, not billed fee. Fee-for-service patients realize the full increase. Ask: "What's my realized increase if 60% of my patients are PPO with fee-schedule contracts?"

"I DON'T KNOW MY OVERHEAD."

Use 65%. Industry benchmark. Your accountant can tell you your actual — but 65% is a fine starting assumption.

SUCCESS SIGNAL

SUCCESS SIGNAL

- You raise fees in January (not "eventually")
- You've briefed your team using the talking points
- The front desk gets zero "did prices go up?" callbacks by February
- Your accountant confirms the profit lift on your Q1 P&L

EOB Decoder

Plain-English EOB decoding, patient-portion math in the right order, and a coverage-book entry for every plan you decode. The most PHI-sensitive skill in the kit — de-ID gate is mechanically enforced.

WHAT IT DOES

EOB Decoder takes a **de-identified** EOB and returns a plain-English line-by-line decode, an accurate patient portion (deductible → coverage % → annual-max, in the right order), a plain-English explanation of any denial or remark code with the next-step fix, a coverage-book entry for that plan (so every future patient on it skips the back-and-forth), and — optionally — a patient-friendly estimate sheet.

PHI-SENSITIVE SKILL

This is the most PHI-sensitive skill in the kit. The de-identification gate is enforced mechanically — the skill halts and refuses to compute if it detects a name, DOB, member ID, or other identifier.

BEST-FIT SCENARIO

- An EOB just came in and the patient balance doesn't match what you collected
- A claim was denied with a code you don't understand
- You're onboarding a new front-desk hire and want them to stop calling the payer for every remark code
- You've never built a coverage book, and you know you should

Frequency: as-needed. Practices that use this on every unusual EOB build a coverage book in ~90 days that saves 10–15 hrs/month of insurance calls thereafter.

INPUTS (WHAT YOU PROVIDE)

Strip BEFORE pasting (the built-in `detect_phi()` runs as a hard gate):

- Patient name
- Member ID / Subscriber ID
- Claim number
- DOB
- Address (patient and subscriber)
- Identifying group name (generic plan name like "PPO Basic 2024" is fine; "ACME Corp Dental" may identify)

Keep and share (safe under HIPAA Safe Harbor):

- CDT procedure codes (D2750, D4341, etc.)
- Tooth numbers
- Billed / allowed / paid / patient-responsibility amounts
- Coverage %, annual max (total and remaining), deductible (met and total)
- In-network or out-of-network status
- Remark codes and reason codes (CO-45, PR-2, etc.)
- Plan type (PPO, DHMO, Indemnity)

You can paste the EOB text OR upload a redacted EOB image/PDF (with identifiers blacked out).

OUTPUTS (WHAT CLAUDE PRODUCES)

A) Line-by-line decode (for the front desk):

```
D2710 (Porcelain crown, tooth 19)
  Billed: $1,350 | Allowed: $850 | Downgraded to resin rate
  Deductible applied: $100 | Plan pays: $600 | Patient owes: $250
  Note: Crown was covered at the lower amalgam/resin rate.
  If you want the porcelain crown in future, a narrative + radiograph
  can support an appeal.
```

B) Coverage-book entry (internal use): a standardized record you file for this plan. Every future patient on it inherits the benefit info. Each entry carries "Verify benefits with payer before treatment. Amounts are estimates."

C) Patient estimate sheet (optional, patient-facing): Plain language, two lines — "Pay in full today: $X" and "Monthly estimate: $X/mo over Y months (estimate, not a guarantee)." No insurance jargon, no claim numbers, no remark codes.

THE EXACT PROMPT TO TRIGGER IT

PROMPT

Run Chairside EOB Decoder. I have a de-identified EOB to decode.

Then paste the EOB text (with all identifiers stripped) OR upload the redacted image/PDF. The skill runs the PHI gate FIRST — before echoing or computing anything.

WORKED EXAMPLE

WORKED EXAMPLE

INPUT (PASTE TO CLAUDE, DE-IDENTIFIED)

Plan: PPO Basic (Delta Dental) · In-network: yes
Annual max: $1,500 total, $1,400 remaining
Deductible: $50 met of $50 total
Coverage: Basic 80% / Major 50%
Lines:

- D2710 (tooth 19) — billed $1,350 / allowed $850 / remark: downgrade to D2140 rate
- D2140 (tooth 3) — billed $215 / allowed $175 / remark: PR-2 (deductible)

LINE-BY-LINE DECODE

D2710 (Porcelain crown, tooth 19)
Billed: $1,350 · Allowed: $850 · Downgraded to resin rate (D2140)
Deductible applied: $0 (deductible was met by the fillings line)
Plan pays: 50% of $850 = $425
Patient owes: $850 – $425 = **$425** (plus write-off portion of $500)
Note: The porcelain crown was covered at the lower alternate-benefit rate. If the patient prefers to appeal for the porcelain rate in the future, a narrative + preoperative radiograph can support the appeal.

D2140 (Composite filling, tooth 3)
Billed: $215 · Allowed: $175
Deductible applied: $50 (met the full deductible on this line)
Plan pays: 80% of ($175 – $50) = $100
Patient owes: $175 – $100 = **$75** (plus write-off of $40)

Total patient responsibility: $500

COVERAGE-BOOK ENTRY — DELTA DENTAL PPO BASIC 2024 (INTERNAL)

- Annual max: $1,500

- Deductible: $50 individual
- Basic: 80% (fillings, extractions, endo)
- Major: 50% (crowns, bridges, endo, perio surgery)
- Alternate-benefit clause: yes (downgrades porcelain crowns to resin rate on posterior)
- Waiting period: none (per this EOB — verify at benefits check)
- Appeal support: narrative + radiograph accepted for porcelain crown appeals

Verify benefits with payer before treatment. Amounts are estimates.

PATIENT ESTIMATE — CASE AT TOOTH 19

Pay in full today: **$500**
Or: **$45/mo over 12 months** with our financing partner (estimate, not a guarantee of approval).
This estimate cannot be verified with payers in real time. Actual balance may differ.

TROUBLESHOOTING

"IT HALTED AND WON'T COMPUTE — SAID IT DETECTED AN IDENTIFIER."

The regex tripped. Look at your paste. Common misses: a phone number in a footer, a claim number, a patient name in the "Insured" field. Scrub and re-paste. Never ask the skill to override the halt.

"THE MATH DOESN'T MATCH WHAT THE PAYER SAYS."

Three common causes:

1. **OON estimate:** the skill halves the allowed amount for OON — conservative by design. Actual payer reimbursement may be higher. Call the payer to reconcile.
2. **Shared-network fee schedule:** some PPOs use shared networks with different schedules than what's on the EOB. Skill can't see those. Verify with payer.
3. **Coverage % came in as 0.8 not 80:** the skill auto-scales but flags it. Read the coverage_note in the output.

"THE REMARK CODE ISN'T IN THE SKILL'S DICTIONARY."

The skill loads a common-codes reference. If it doesn't recognize a code, it'll say so and ask what it means. Give it the payer's official definition and it'll incorporate that.

"I NEED TO APPEAL A DENIAL."

The skill doesn't write appeals directly (that's clinical judgment). But it will tell you what a good appeal for that specific denial code requires — narrative? Radiograph? Peer review request? Then you write the appeal, and use Claude generally (not this skill) to help polish the language.

"CAN IT READ AN EOB IMAGE?"

Yes — upload the redacted image or PDF. It OCRs and runs the PHI gate on the OCR output. Still redact before you upload; OCR might fail on a black rectangle over a name.

SUCCESS SIGNAL

SUCCESS SIGNAL

- Your front desk stops calling payers for routine EOB questions
- Your coverage book has 15+ plans in it after 90 days
- The math your team collects at check-out matches what the EOB says, ± a few dollars
- Denials get faster resolutions because the fix is already in the decode

Cyber Self-Audit

18-question guided audit across 8 risk domains. Risk-tier badge, prioritized roadmap, and two copy-paste call-sheets (one for IT, one for insurer). No patient data required.

WHAT IT DOES

Cyber Self-Audit walks you through an 18-question guided audit across 8 risk domains — server age, backup verification, cyber insurance, MFA coverage, IT partnership quality, OS update hygiene, phishing training, and vendor BAAs. You answer; the skill scores. It returns a risk-tier badge (Critical / High / Moderate / Low-watch), an 8-domain traffic-light dashboard, a prioritized roadmap (This Week / This Month / This Quarter), and two copy-paste call-sheets — one for your IT provider, one for your insurer.

No patient data required or accepted. Every input is about your systems and policies.

BEST-FIT SCENARIO

- You've never had a formal cyber audit and want a baseline
- Your cyber insurance is up for renewal and you need to know what's actually covered
- You just heard about a peer practice that got ransomwared and want to know your exposure
- You're switching IT providers and want to compare "before" and "after"

Frequency: annually. Do a fresh score every 12 months. Also rerun after any major infrastructure change (new PMS, new server, new IT provider).

INPUTS (WHAT YOU PROVIDE)

Just answers to 18 questions. Takes 10–15 minutes. **"I don't know" is a valid and important answer** — the skill flags it as a scored gap, not "probably fine."

Domain by domain:

1. **Server & Backup:** server age (years), offsite backup yes/no
2. **Verified Cloud Restore:** IT ever done a live restore test, if yes how many months ago
3. **Cyber Insurance:** standalone policy yes/no, ransomware covered, policy limit adequate for 2–4 weeks downtime

4. **MFA Coverage:** MFA on email, remote access, PMS
5. **IT Partnership Quality:** IT sends regular security reports, patches proactively or reactively
6. **OS & Update Hygiene:** all OSs up to date, auto-updates on or manual
7. **Phishing Training:** team ever done phishing simulation, how often it repeats
8. **Vendor BAAs:** signed BAAs with vendors handling patient data, do you know how many

Optional: upload your cyber insurance declarations page (policy # blacked out) and the skill will read it and flag ransomware exclusions, sublimits, and retroactive-date gaps.

OUTPUTS (WHAT CLAUDE PRODUCES)

- **Risk-tier badge** — Critical / High / Moderate / Low-watch — with threshold explanation
- **8-domain traffic-light dashboard** — Red / Yellow / Green / Verify for each domain
- **Red-flag callouts** — any condition that auto-caps your tier (e.g., "no offsite backup" is an automatic Critical)
- **Prioritized roadmap:** This Week (1–3 moves, cheap first step, who owns it) · This Month (30-day improvements) · This Quarter (longer-horizon)
- **Two copy-paste call-sheets:** IT questions (with "a good answer sounds like...") + Insurer questions (same format)
- **Insurance policy notes** (if declarations page uploaded) — flags ransomware exclusion, sublimit, retroactive-date gap, labeled "review with your broker"

THE EXACT PROMPT TO TRIGGER IT

> PROMPT
>
> *Run Chairside Cyber Self-Audit.*

The skill will introduce itself, note the "I don't know" rule, and start with Domain 1. Answer one domain at a time. Total time ~12 minutes.

WORKED EXAMPLE

> WORKED EXAMPLE
>
> **Practice:** solo GP, 4 ops, one server in the practice.

SAMPLE EXCHANGE (EXCERPT)

Claude: "Domain 1 of 8 — Server & Backup. Q1: How old is your main practice server?"
You: "About 7 years."
Claude: "Q2: Does your practice have an offsite or cloud backup — a copy of your data that lives somewhere other than the office?"
You: "IDK — I think so? IT set something up last year."
Claude: "Noted — I'll flag that as 'verify.' Domain 2 of 8 — Verified Cloud Restore. Q3: Has your IT provider ever done a live restore test..."

(...continues through all 8 domains...)

FINAL OUTPUT (EXCERPTS)

PRACTICE CYBER RISK SELF-AUDIT — [date]
OVERALL RISK TIER: HIGH (score 42/100)

RED FLAGS:

- Server age 7+ years (industry benchmark: replace at 5)
- No verified restore test on your cloud backup (an unverified backup is not a backup)
- Cyber insurance status unknown (call your broker this week)

DOMAIN DASHBOARD

Server & Backup: **Yellow (42/100)** — Server old; backup exists but unverified

Cloud Restore: **Red (20/100)** — Never tested

Cyber Insurance: **Yellow (verify)** — Status unknown; call broker

MFA Coverage: **Green (60/100)** — MFA on email; add PMS

IT Partnership: **Red (18/100)** — Reactive-only IT

OS & Updates: **Yellow (42/100)** — Manual updates

Phishing Training: **Green (60/100)** — Annual training in place

Vendor BAAs: **Green (75/100)** — Most BAAs signed; know your count

ROADMAP

This week:

- [IT] Have your IT provider run a live restore test from your cloud backup. Cheap first step: email today "I need proof of a live restore test within 7 days." Cost: $0.
- [You] Call your broker: "Do we have a standalone cyber liability policy? Does it cover ransomware — both ransom payments and data recovery?" Cost: 10 min.

This month:

- [IT] Enable auto-updates on all workstations and the server.
- [IT] Add MFA on your PMS if the software supports it.

This quarter:

- [IT] Server replacement plan (7 years is well past comfortable lifespan).
- [IT] Shift to a proactive-managed-services IT contract if current provider is reactive-only.

TROUBLESHOOTING

"I ANSWERED 'IDK' TO A LOT. IS THE SCORE UNRELIABLE?"

The opposite. "IDK" is scored as a gap — the skill assumes you can't verify the control, so it's not credit you get. The score is conservatively accurate. Convert IDKs to Yes or No over the next 30–60 days using the call-sheets.

"THE SKILL SAID WE'RE 'CRITICAL' BUT I FEEL FINE."

Any Critical rating has a specific red-flag reason attached. Read it. Common ones: no offsite backup, cyber insurance excludes ransomware, no MFA anywhere. If the red flag is real, you're not fine — you're lucky nothing has happened yet.

"THE IT CALL-SHEET USES TERMS I DON'T UNDERSTAND."

That's fine. The call-sheet gives you the question AND "a good answer sounds like..." — so you know what a competent answer looks like. If IT gives you a vague, jargon-laden non-answer to a specific question, that IS the signal.

"I DON'T HAVE CYBER INSURANCE AT ALL."

Cyber Self-Audit will still score you — you'll get a red flag on Domain 3. Use the insurer call-sheet to shop for a quote. Ballpark for a small practice: $1,500–$3,000/yr for adequate coverage.

"SHOULD I SHARE THIS WITH MY IT PROVIDER?"

Yes — that's the point. Print it, hand it to them, ask them to help you close the Reds by next quarter. If they resist or dismiss it, you have a partnership-quality problem (Domain 5), not a technical problem.

SUCCESS SIGNAL

SUCCESS SIGNAL

- You have real answers (not IDK) to every question within 60 days
- Your This Week actions are done within 7 days
- You brought this to your IT provider and they took it seriously
- Next year's score improves by at least one tier

Chapter 09 · Skill 6 of 8

Morning Huddle

A one-page printable huddle sheet from your de-identified day-sheet. Route-to-closer flags for high-value cases sitting in hygiene, open-chair fill plan, one team talking point. Under a minute to run.

WHAT IT DOES

Morning Huddle takes your **de-identified** day-sheet (initials only, no names) and returns a one-page printable huddle sheet — today's scheduled production vs. daily goal, split by provider, high-value case routing flags (which HV case is sitting in a hygiene slot and needs the TC), open-chair fill plan, unscheduled treatment coming in today, and one team talking point.

Runs in under a minute. Print it. Hang it in the huddle room. Run the 10-minute huddle.

BEST-FIT SCENARIO

- Your huddles are chaotic or you skip them because they take too long to prep
- You have a strong TC (or want to develop one) and need a system to route HV cases to her instead of the hygienist
- Your daily production varies wildly and nobody's tracking to the monthly goal in real-time
- You want a printable one-pager the whole team looks at, not a whiteboard scribble

Frequency: every practice day. 3 minutes each morning. **The daily-driver skill of the kit.**

INPUTS (WHAT YOU PROVIDE)

Strip BEFORE pasting:

- Full patient names (use initials only — "R.K." not "Robert Kelley")
- DOB, phone, address, insurance ID

Keep and share: Time · Provider/Op · Initials · Appointment type · Scheduled $ · Tx-plan status

Format: one appointment per line, pipe-delimited:

```
8:00 AM | Dr. Patel / Op 1 | R.K. | New Patient Exam + X-rays | $285 | needs tx
plan
8:00 AM | Hygiene / Op 3   | M.T. | Prophy + Implant Consult  | $8,400 | tx
planned
```

```
9:00 AM | Dr. Patel / Op 1 | S.L. | Crown Prep (PFM)          | $1,450 | tx
planned
```

Mark open-chair blocks with a comment line: `# --- Op 2 open 10:00 AM – 12:00 PM ---`

One-time Project setup (recommended): Create a Claude Project called `[Your Practice] – Huddle` and add a `practice-config.md` with practice_name, providers, monthly_goal, working_days, hv_threshold. Every morning: open the project → paste today's schedule → run → get sheet.

OUTPUTS (WHAT CLAUDE PRODUCES)

SECTION	WHAT IT SHOWS
Today's number	Scheduled production / daily goal / % to goal / $ gap
By provider	Subtotal per provider or operatory
High-value money list	Each HV case: initials, type, $, who closes, next action
Route-to-closer flags	HV cases in hygiene slots — explicit TC assignment
Unscheduled treatment	Patients with open tx plans coming in today
Open-chair fill plan	Which blocks are open + fill priority
Team talking point	One rotating topic (clinical / culture / goal)

THE EXACT PROMPT TO TRIGGER IT

PROMPT

Run Chairside Morning Huddle. Daily goal $6,000. Here's today's schedule: [paste]

WORKED EXAMPLE

WORKED EXAMPLE

INPUT (PASTE TO CLAUDE)

Daily goal: $6,000. Practice: Sunrise Family Dental.

```
8:00 AM  | Dr. Patel / Op 1  | R.K.  | New Patient Exam + X-rays  | $285   | needs tx plan
8:00 AM  | Hygiene / Op 3    | M.T.  | Prophy + Implant Consult   | $8,400 | tx planned –
full arch
9:00 AM  | Dr. Patel / Op 1  | S.L.  | Crown Prep (PFM)           | $1,450 | tx planned
10:00 AM | Dr. Patel / Op 1  | B.G.  | Composite ×2               | $520   | tx planned
10:30 AM | Hygiene / Op 3    | J.R.  | Perio Maint                | $185   | –
11:00 AM | Dr. Patel / Op 1  | T.P.  | Emergency Exam             | $180   | needs tx plan
# --- Op 2 open 10:00 AM – 12:00 PM ---
```

```
1:00 PM  | Dr. Patel / Op 1  | K.M.  | Extraction #14            | $450   | tx planned
2:00 PM  | Dr. Patel / Op 1  | L.H.  | Impl restoration          | $2,400 | tx planned
# --- Op 1 open 3:00 – 4:00 PM ---
```

Claude echoes: "I read 8 appointments totaling $13,870 — correct? Before I generate the huddle sheet, confirm this matches your schedule."

CLAUDE PRODUCES

MORNING HUDDLE — [date] · Sunrise Family Dental

Today's number: Scheduled **$13,870** · Daily goal $6,000 · **231% to goal** · Over by $7,870

By provider: Dr. Patel: $5,285 (6 appts) · Hygiene: $8,585 (2 appts)

High-value money list:

- M.T. — Full Arch Implant Consult **$8,400** — sitting in hygiene → **route to TC before patient leaves the chair**
- L.H. — Implant restoration **$2,400** — Dr. Patel closes

Route-to-closer flags: M.T. at 8:00 AM in hygiene. Assign TC to arrive at 8:45 for warm handoff before the prophy ends. Consult should not leave without booking.

Unscheduled treatment coming in today: R.K. at 8:00 AM — needs tx plan built (Dr. Patel to present) · T.P. at 11:00 AM — emergency exam, likely restorative

Open-chair fill plan: Op 2 open 10 AM – 12 PM: call ASAP list first, then reactivation · Op 1 open 3–4 PM: fill from same-day + ASAP list

Team talking point today: You're on pace to hit monthly goal by day 16 if today's book holds. Take the win in huddle. Recognize whoever booked M.T.'s consult.

TROUBLESHOOTING

"THE ECHO-BACK SAYS A DIFFERENT TOTAL THAN MY PMS."

Two common causes: PMS may include holds or "TBD" appointments in the total (verify against actual scheduled). Or you included a bundled fee that's really the second visit's fee (adjust the paste to show only today's portion).

"IT DIDN'T FLAG A HV CASE I KNOW IS COMING."

Check the $ amount — the default HV threshold is $3,000. If your case is below, it won't flag. Adjust in your practice-config.md: `hv_threshold: 2000`.

"THE ROUTE-TO-CLOSER FLAG MAKES MY HYGIENIST MAD."

Reframe: the flag isn't a criticism of hygiene — it's protection. Hygienists have 30-minute slots and a mouthful of instruments. A $8K consult needs 60 minutes and a closer's full attention. Bring the flag to your team meeting; explain the logic.

"THE TALKING POINT FEELS GENERIC."

The skill rotates through clinical / culture / goal. Ask: "Give me three alternative talking points for today — one clinical, one culture, one goal." Pick your favorite.

"I CAN'T DO THIS EVERY MORNING — I DON'T HAVE TIME."

It's 3 minutes. If your practice can't spare 3 minutes for a huddle prep, you have a scheduling problem, not a skill problem. Do it while your coffee brews.

SUCCESS SIGNAL

SUCCESS SIGNAL

- Your huddles are shorter, sharper, and hit the same beats every day
- Your HV consults are getting closed the same day, not 5 days later
- Your team knows the number by 8:15 AM every day
- You're catching open-chair blocks BEFORE the chair is empty, not after

Hiring Kit

A complete, EEO-clean hiring packet in one Artifact: job post, phone-screen questions, structured interview scorecard, paid working interview, reference checks, 30/60/90 onboarding, decision summary.

WHAT IT DOES

Hiring Kit takes a role (hygienist / TC / front desk / DA / associate) and your practice's 3–5 stated values, then produces a complete, EEO-clean hiring packet in one Artifact: job post, 5–7 phone-screen questions, structured interview scorecard with anchored 1–5 ratings, paid working interview task with rubric, reference-check questions, 30/60/90 onboarding checklist, and a one-page decision summary.

Every section is EEO-compliant. No blocked questions. No demographic probes. Compensation shows only figures you provided — the skill never invents a number.

No patient data ever needed. Applicant information (résumé, notes) is employment data, not PHI — normal candidate confidentiality applies.

BEST-FIT SCENARIO

- You're hiring for the first time in 12+ months and your job posts are stale
- You've had bad hires and you know the interview process is the weak link
- You just lost someone and you're hiring under time pressure — the packet gives you structure fast
- You want your team to hire without you having to write every interview question

Frequency: as-needed. Every open role should get a fresh packet — values evolve, comp changes, must-haves shift.

INPUTS (WHAT YOU PROVIDE)

- **Role:** hygienist / TC / front desk / DA / associate (or a hybrid)
- **3–5 practice values** — what your practice genuinely stands for. If unsure, paste this to your team chat: *"Name three things we do here that you'd miss if you left."* Values come from the answers, not invention.

- **Compensation:** hourly, salary, production %, or combination. Use `[owner to set]` for figures you're not ready to commit — the skill will never invent one.
- **Must-haves:** non-negotiables (license, years of experience, bilingual, etc.)
- **Nice-to-haves:** preferred but not disqualifying
- **Context:** solo vs. group, insurance mix, growth phase, why the last person left
- *Optional:* paste an applicant's résumé text — the skill will generate a live scorecard map for that candidate

OUTPUTS (WHAT CLAUDE PRODUCES)

Seven sections in one Artifact:

1. **Job post** — EEO-clean, values-grounded, no hyperbole, compensation shown or `[owner to set]`, required EEO statement
2. **Phone-screen questions** — 5–7 structured questions: logistical, behavioral, culture signal. No blocked probes.
3. **Structured interview scorecard** — 5–7 competencies from role + values. Each: behavioral question + 1–5 anchored ratings. Total score + gut-check free-text.
4. **Paid working interview task** — 15–60 min, role-appropriate, PHI-free. Rubric with 3–5 observable behaviors, each rated 1–3. Fairness note.
5. **Reference-check questions** — five targeted for past supervisors. Each with "strong answer looks like" and "red flag looks like."
6. **30/60/90 onboarding checklist** — three milestone tables. Clinical immersion first, role-specific milestones, owner/buddy checkpoints, success metric per phase.
7. **One-page decision summary** — printable, filled after interviews. Scorecard total, working interview total, reference pass/flag, offer/hold/decline decision.

THE EXACT PROMPT TO TRIGGER IT

> PROMPT
>
> *Run Chairside Hiring Kit. I'm hiring a treatment coordinator. Our values are: honest with patients, own your outcomes, we grow together. Comp is $22–26/hr base plus 2% of net production she closes, capped at $500/mo. Must-haves: 2+ years dental TC or high-end sales, comfortable with financing conversations. Nice-to-have: bilingual English/Spanish. Solo practice, fee-for-service, transitioning from front-desk-doubles-as-TC to dedicated role.*

WORKED EXAMPLE

WORKED EXAMPLE

SECTION 1 — JOB POST (PARTIAL)

Treatment Coordinator — Sunrise Family Dental

We're a fee-for-service solo practice looking for our first dedicated treatment coordinator. Right now our front desk does this on top of everything else, and both roles suffer for it. If you love helping patients understand and choose care — and if you've watched a good financing conversation open a door that was closed a minute ago — we want to meet you.

What you'll actually do:

- Present treatment plans and financial options to patients after the doctor's exam
- Own the follow-up on unaccepted plans — email, phone, warm handoffs
- Coordinate CareCredit and Sunbit applications for patients who need financing
- Track case-acceptance rate and close rate — you'll own the number

Compensation: $22–26/hr base + 2% of net production you close (capped at $500/mo). Base within range depends on experience.

Sunrise Family Dental is an Equal Opportunity Employer. We do not discriminate on the basis of race, color, religion, sex, sexual orientation, gender identity, national origin, age, disability, genetic information, or veteran status under federal, state, or local law.

SECTION 3 — INTERVIEW SCORECARD (EXCERPT)

Competency: Financing Conversation Comfort (Value tie: honest with patients)
Behavioral question: "Tell me about a time you had to help someone make a decision about money — a purchase they wanted but weren't sure they could afford. What did you say? What happened?"

Anchors:

1 = Deflects; can't recall a specific example; frames money as "hard to talk about"
3 = Recalls a specific example; matter-of-fact tone about money; treated the person as capable of deciding
5 = Multiple examples; describes a specific technique (e.g., "I always show the total AND the monthly, so they see both"); talks about money without discomfort

SECTION 6 — 30/60/90 ONBOARDING (DAYS 1–30, EXCERPT)

Days 1–30 — Clinical Immersion + Shadow

- Week 1: Shadow every provider (doctor, hygienist, DA) — one day each. No case presenting yet.

- Week 1: Complete CareCredit and Sunbit training modules.
- Week 2: Sit in on every doctor's exam. Take notes on the "why" behind treatment recommendations.
- Week 3: Present under the doctor's supervision — doctor introduces, TC presents, doctor closes.
- Week 4: Solo presentation with post-review by owner. Review recordings if available.
- **Success metric for Day 30:** TC has presented 8+ plans, understands financing options, has a relationship with at least one patient who chose treatment based on her presentation.

TROUBLESHOOTING

"THE JOB POST FEELS TOO CORPORATE."

It shouldn't. If it does, the skill missed your voice. Say: "Rewrite the job post in a warmer, more conversational tone — the way I'd tell a friend about the role." It'll adjust.

"AN INTERVIEW QUESTION FEELS DISCRIMINATORY."

Send it back: "Question 3 feels like it's probing marital status. Rewrite it." The skill runs an EEO pass but nothing catches everything — you're the human in the loop.

"I DON'T HAVE VALUES WRITTEN DOWN."

The skill will help you extract them. Say: "Interview me for 3 minutes on my practice values before writing the packet." It'll ask you the "three things you'd miss" question and use your answers.

"THE COMP SHOWS `[OWNER TO SET]` IN A LOT OF PLACES."

Fill them in when you're ready. The skill deliberately refuses to invent numbers because bad comp language costs practices. Come back to the packet after you've talked to your accountant about market rate.

"HOW DO I USE THE RÉSUMÉ-SCORING FEATURE?"

Paste the résumé text (copy/paste from PDF is fine). The skill runs a candidate-privacy note, then maps résumé claims against the scorecard. It won't rank candidates for you — that's your decision — but it'll tell you where to probe in the interview.

SUCCESS SIGNAL

SUCCESS SIGNAL

- Your job posts get more (and better) applicants
- Your interviews are shorter and more structured (no "so tell me about yourself")
- Your 90-day retention improves (structured onboarding does this)
- The person you didn't hire tells a friend the interview was fair — even though they didn't get the job

Vision Coach

A coaching tool, not a numbers tool. Describe a people situation; the skill diagnoses, produces a word-for-word script, and offers to role-play the team member so you can rehearse.

WHAT IT DOES

Vision Coach is a coaching tool, not a numbers tool. Describe a people situation — a hard team conversation, a leadership block, a vision-vacuum moment — and the skill will:

1. Hold you to behavior-language instead of labels
2. Run an escalation guardrail to catch anything HR- or legal-adjacent
3. Produce a Socratic reframe of what's actually happening under the surface
4. Once you confirm the reframe, hand you a word-for-word conversation script, a 15-minute meeting agenda, anticipated pushback with replies, and a follow-up checkpoint
5. Offer to role-play the team member so you can rehearse before you walk into the room

WHAT IT WILL NOT DO

- Calculate bonuses, payroll, or KPIs
- Draft termination language or stated employment consequences
- Handle discrimination, harassment, wage/hour, or mental-health-crisis situations
- Give legal, HR, or mental-health advice

For any of those, it stops and refers you to an employment attorney, your state dental association's HR hotline, or a licensed professional.

BEST-FIT SCENARIO

- You've been dreading a hard conversation and keep putting it off
- Your team feels off but you can't name why, and you need help diagnosing
- You have to cast a vision to your team and don't have the words yet

- You've had the same conversation three times and it isn't sticking — the reframe is probably needed

Frequency: as-needed. Some months, weekly. Some months, never.

INPUTS (WHAT YOU PROVIDE)

Guided intake — the skill asks these one at a time:

1. **The situation.** Behaviors, specific, recent. ("She's missed the morning huddle twice this week" — NOT "She's lazy.")
2. **The goal.** What does success look like? What changes, and how do you know?
3. **What's been tried.** What have you already done or said?
4. **Team context.** Role and first name only. No last names, no DOBs. Any relevant tenure/dynamics.
5. **Owner's self-read.** 1–10, how confident are you that YOU are not part of this problem? (Honest answer only.)
6. **The stakes.** What happens if this stays the same in 90 days?

No patient data. Ever. No employee identifiers beyond first-name/role.

OUTPUTS (WHAT CLAUDE PRODUCES)

A one-page Artifact:

1. **The reframe** — one short paragraph: what's really happening, named plainly
2. **Leadership framing** — manage-vs-lead distinction applied to this situation
3. **Word-for-word opening script** — opening line, 3–5 behavioral asks, the ask (framed through patient care not revenue), the close (concrete, time-bound)
4. **Meeting or huddle agenda** — a tight 15-minute agenda with time blocks
5. **Anticipated pushback + replies** — 3–4 most likely, with word-for-word responses
6. **Follow-up checkpoint** — concrete: what you'll observe, when, what you'll say if nothing has changed
7. **Vision statement** (only when the reframe is a vision vacuum)

Then it offers to role-play the team member so you can rehearse.

THE EXACT PROMPT TO TRIGGER IT

> **PROMPT**

> *Run Chairside Vision Coach. I need to have a hard conversation with our lead hygienist.*

The skill starts the guided intake. Answer one question at a time.

WORKED EXAMPLE

WORKED EXAMPLE

THE SITUATION (OWNER'S PASTE)

Lead hygienist, "Sarah," 8 years with us. In the last 6 weeks, she's stopped presenting perio treatment plans to patients who need them. I've watched her diagnose D4341-level cases and then just do a prophy anyway. Our perio case acceptance has dropped 40%.

Claude's intake continues through the six questions...

ESCALATION SCREEN

No termination language, no protected class, no wage/hour. Proceeds.

BEHAVIOR-LANGUAGE CHECK

Owner used "she's stopped" (behavior), not "she's checked out" (label). Passes.

SOCRATIC REFRAME

Here's what I hear under the surface:

Sarah has been with you 8 years — she's the most senior clinician on your team after you. Six weeks ago something changed. You said in the intake you scored yourself a 4 on "am I part of this problem." That's an honest answer, and it points at where I want to go.

This looks less like a compliance problem and more like a **leadership-vacuum problem**. Sarah made a decision six weeks ago to stop presenting perio. She had a reason. You don't know what it is yet — because you've been managing the outcome (case acceptance dropped) instead of leading the conversation (what happened six weeks ago?).

The management move is to tell her to present perio plans. That'll work for a week. The leadership move is to sit down and ask her what happened.

Does that reframe land, or am I missing something?

WORD-FOR-WORD OPENING SCRIPT (EXCERPT)

Opening line: "Sarah, I want to talk about something. It's not a performance issue. I've noticed something changed about six weeks ago in how you handle perio patients, and I want to understand it. Can we take 30 minutes tomorrow morning?"

Big rocks:

1. In the last six weeks, our perio case acceptance dropped 40%. I've watched you diagnose cases that meet 4341 criteria and then do a prophy anyway.
2. I trust your clinical judgment. This is why I want to understand what changed — not tell you what to do.
3. If there's something about the referral process, the treatment plan presentation, or a patient interaction that made you back off, I want to hear it.
4. If it's something about me or our team's culture, I want to hear that too.

The ask: I need our perio patients to get the care they need. That's why we're both here. What would help you feel confident presenting these plans again?

The close: Whatever comes out of this conversation, let's decide together on one thing you'll try over the next 30 days, and let's check back in on Friday two weeks from now.

ANTICIPATED PUSHBACK + REPLY (EXCERPT)

Pushback: "The patients don't want it. They can't afford it."
Reply: "Some can't. Most can — we have financing partners for a reason. I'm hearing you might have run into a stretch of patients who couldn't, and that made this feel harder. Is that close?"

Role-play offer: "Want me to play Sarah? I'll respond as she probably would, and you can practice your opening line."

TROUBLESHOOTING

"THE REFRAME FEELS OFF — I DON'T AGREE."

Correct it. Say: "The reframe misses X — actually she just came back from maternity leave and I think there's a confidence issue." Claude adjusts and rewrites everything downstream.

"THE SCRIPT IS TOO SOFT."

Ask: "Make the ask firmer — she needs to understand this can't continue at the current level." Just don't ask it to state an employment consequence. That's where the skill stops.

"IS THIS HR TERRITORY OR COACHING?"

The skill's escalation guardrail runs at intake — if it detects HR/legal risk, it stops and refers you out. If it proceeds, you're in coaching territory. Trust the gate; it's calibrated conservatively.

"I WANT THE SCRIPT FRAMED AROUND REVENUE."

The skill won't do that in a 1:1 behavioral conversation. Reason: framing behavior conversations around revenue reads as pressure and produces short-term compliance, not durable change. Have the revenue conversation in a separate all-hands context.

"CAN IT HELP WITH VISION-CASTING TO THE WHOLE TEAM?"

Yes — say "This is a vision-casting session for a team retreat, not a 1:1." It shifts into vision-statement mode and uses a different framework (the blind-five-things exercise, if you haven't got a vision statement yet).

SUCCESS SIGNAL

SUCCESS SIGNAL

- You have the hard conversation instead of putting it off another week
- The conversation actually changes the behavior — you observe the change at the follow-up checkpoint
- You've caught yourself using behavior-language instead of labels when talking about your team
- You've role-played at least one opening line before walking into a hard conversation

Chapter 12

The 90-day rollout plan

Do this in order. Don't skip. Don't front-load. Follow the plan or invent your own — but not both.

WEEK 1

DAYS 1–7

Money + Case

INSTALL · MONEY-LEAK FINDER · CASE CLOSER

Run:

- **Monday:** Export 60–90 days of bank statement, run Money-Leak Finder. Make 3 calls by end of week.
- **Tuesday–Friday:** Run Case Closer on every plan >$3K. Build muscle memory.

Why this pair: Highest immediate ROI (Money-Leak Finder finds real money in week 1) + highest recurring lift (Case Closer improves acceptance on every plan going forward). If you only run 2 skills ever, run these.

Success signal by Friday: You've cut $300+/mo of confirmed waste AND you've handed at least 3 patients a Case Closer handout.

WEEKS 2–4

DAYS 8–28

Fee + Huddle

INSTALL · FEE AUDITOR · MORNING HUDDLE

Run:

- **Week 2:** Run Fee Auditor once. Take the recommendations to your accountant + your team.
- **Weeks 3–4:** Run Morning Huddle every practice day. Set up the Project with your practice config so it's <3 min each morning.

Why this pair: Fee Auditor is a once-a-year action — but the once matters (the profit lift is real). Morning Huddle is daily — but building the daily habit takes 3 weeks, so start it now while momentum is high.

Success signal by end of Week 4: You have a fee-increase plan for January (or now, if you're bold) AND your team is looking at a printed huddle sheet every morning.

MONTH 2

DAYS 29–60

EOB + Hiring (as needed)

INSTALL · EOB DECODER · HIRING KIT

Run:

- **EOB Decoder:** as denials and unclear EOBs come in. Build the coverage-book habit — every plan you decode adds to the book.
- **Hiring Kit:** only if you're hiring. If not, install it and skip it. Come back when there's a role open.

Why this pair: Both are situational, not scheduled. Install them now so you don't have to remember to install them when the moment hits.

Success signal by end of Month 2: Your coverage book has 5–10 plans in it. If you're hiring, you have a packet, not a scramble.

MONTH 3

DAYS 61–90

Cyber + Vision (annual/quarterly cadence)

INSTALL · CYBER SELF-AUDIT · VISION COACH

Run:

- **Cyber Self-Audit:** once. Get your score. Take the roadmap to your IT provider and insurer. Rerun in 12 months.
- **Vision Coach:** as needed. Most owners use it 1–2x per quarter — the hard conversations don't happen every week, but they happen.

Why this pair: Both are lower-frequency, higher-stakes. You want them installed and familiar BEFORE the crisis moment — the ransomware scare, the hard conversation you've been dreading. Not during.

Success signal by end of Month 3: You've done your annual cyber audit AND you've had at least one hard conversation you'd been dreading.

After Day 90 — the ongoing cadence

By Day 90, all 8 skills are installed and you've used all of them at least once. From here:

- **Money-Leak Finder** — every 3 months
- **Case Closer** — 2–5x per week
- **Fee Auditor** — once per year
- **EOB Decoder** — as needed (every unusual EOB)
- **Cyber Self-Audit** — once per year (unless infrastructure changes)
- **Morning Huddle** — every practice day
- **Hiring Kit** — every open role
- **Vision Coach** — as needed

This is now how you work. Not a project. A default.

Chapter 13

What NOT to use these skills for

The kit is powerful because it stays in its lane. Do not put it in the wrong lane.

DO NOT — Any workflow that requires PHI

- Clinical notes with patient identifiers
- Treatment plans with a specific patient's full name attached
- Insurance claims tied to a specific patient
- Lab orders, prescriptions, patient-communication drafts that include the patient's name

Why: Consumer Claude and consumer ChatGPT are NOT HIPAA-compliant. There is no Business Associate Agreement (BAA) in place. Zach Hanson's session covered the correct pattern.

The rule: Strip patient identifiers BEFORE the prompt → send de-identified content → get de-identified output → human review → re-attach patient context manually on the practice side.

Every patient-adjacent skill in this kit (EOB Decoder, Case Closer, Morning Huddle) has a hard-coded de-identification gate. That gate is a backstop, not a guarantee. YOU strip the identifiers.

If your workflow can't be de-identified: you need a BAA-backed AI vendor, not this kit. Zach's session listed the current options.

DO NOT — Clinical decision support

- "Should I do a crown or an onlay?"
- "What's the differential for this radiolucency?"
- "Prescription dose for a 200-lb male with penicillin allergy?"

Why: These skills are not FDA-cleared clinical decision-support tools. They're operational and financial tools. Ask a colleague, a specialist, or a clinical decision-support product built for this purpose.

DO NOT — Specific-patient cost estimates with real-time insurance math

Example: "What will Mrs. Jones owe on this crown?"

Why: Real-time benefit verification requires a payer connection. EOB Decoder decodes an EOB after the fact — it does not verify benefits in real time. For live benefit verification, use your practice management software or a benefit-verification service (Zuub, Vyne, etc.).

DO NOT — Legal, HR, or tax advice

- "Can I fire this person?"
- "Am I compliant with OSHA?"
- "What's the tax treatment of my new equipment purchase?"

Why: These skills are coaching and operational tools. For legal/HR/tax, call the appropriate professional. Vision Coach explicitly stops the moment a conversation crosses into HR/legal territory — designed to refuse, not to try.

DO NOT — Anything you would not want shown in a deposition

If you'd be embarrassed to have the paste and the output shown to your patients, staff, or malpractice insurer — don't paste it. This is the sniff test.

THE LANE TEST

Before you paste anything into any of these skills, ask:

1. **Is this operational or clinical?** Operational → maybe. Clinical → no.
2. **Is this PHI-free at the paste?** Yes → maybe. No → strip it first.
3. **Am I asking for a draft I'll review, or a decision I'll act on?** Draft → yes. Decision → no.

If all three are green, run the skill.
If any is yellow, pause.
If any is red, don't paste.

Built for the Bulletproof Dental CE Summit by AIPG — Blake McClellan.
Kit source: `github.com/aipg-os/chairside-ai` *(or check the resource page for the current download).*
Questions: [Blake's email TBD] · [Resource page URL TBD]

END OF GUIDE

Total install time: 15 minutes for one skill · 45 minutes for all 8. Total time to first return: 5 minutes.