


FROM DEMO TO PRODUCTION

Claude in *Healthcare*.

Three deployment patterns for AI in healthcare — from Anthropic's native connectors, through RCM automation, to a single-source-of-truth data warehouse inside a diagnostic lab.

DIAGRAM 00 · ONE SAFE POINT OF CONTACT



6 min → 30
S

Time to process a single denial after RCM automation deployment.

80–85%

Share of BannerWise users reporting clinical time savings.

14 hrs/wk

Time prior authorization steals from the average U.S. physician.

What's in this ebook.

Three deployment patterns. One real case study. Six pitfalls. A 90-day checklist. No hype, with numbers from production.

01	Introduction	p. 04
	Why 2026 is the year of production AI in healthcare.	

02	Claude for Healthcare	p. 06
	What Anthropic actually shipped in January 2026.	

03	The compliance foundation	p. 09
	HIPAA, BAA, Zero Data Retention — what actually protects you.	

04	Three deployment patterns	p. 11
	Native Connectors · RPA + Claude · Data Warehouse.	

05	Case study: a diagnostic laboratory	p. 17
	Claude Enterprise, the warehouse, zero MCPs to source systems.	

06	ROI and pattern selection	p. 20
	The 2x2 decision matrix. Where the money actually is.	

07	Six pitfalls	p. 22
	The first deployments that most often cost twice.	

08	The first 90 days	p. 24
	Checklist for days 0–30 · 31–60 · 61–90.	

Karl Mielnicki, Co-founder & CTO at Flobotics. 45 minutes. A written recommendation.

— INTRODUCTION

2026 is not the year of more announcements. It's the year of production.

In January 2026, at the J.P. Morgan Healthcare conference, Anthropic announced **Claude for Healthcare** — the first version of the model purpose-built around concrete, clinically and administratively regulated workflows. Over the next three months, Amazon, Microsoft, and OpenAI added their own announcements.

Healthcare stopped being the industry of "last AI beta-testers" and became the arena where the four largest AI platforms compete for the same dollars: the **\$360 billion wasted annually on administration in the U.S. healthcare system.**

This ebook is not about announcements. It's about *what Claude actually looks like once you plug it into a real organization* — with real PHI, real audit logs, and a real budget.

We're laying out **three deployment patterns** we see across organizations in 2026, from diagnostic labs to large provider networks:

1. **Pattern A — Native Connectors.** Claude for Healthcare out of the box: CMS, ICD-10, PubMed, HealthEx.
2. **Pattern B — RCM process automation.** Claude wired into RPA pipelines, no replatforming. This is the model Flobotics deploys for payers and providers.
3. **Pattern C — Data warehouse as the single source of truth.** The pattern a large diagnostic lab just deployed on Claude Enterprise — *with zero MCPs to source systems.*

WHO THIS EBOOK WILL SURPRISE

Leaders who assume "AI deployment" equals "plugging the model into the EHR."

In reality, the most expensive and most secure deployment of 2026 — described in chapter 5 — **has no integration to any source system.**

DIAGRAM 01 · COMPETITIVE LANDSCAPE

Q1 2026



Three months changed the landscape. The four largest AI platforms in a single quarter: front-office (Microsoft), reasoning (Anthropic), back-end (Amazon, OpenAI).

— CHAPTER 01 • CLAUDE FOR HEALTHCARE

Anthropic didn't ship a new model. They shipped a *product layer*.

Claude for Healthcare isn't "a new Claude." It's a **product layer on top of the existing model**, made of five components. Each one unblocks something that was holding back previous deployments.

1 • CMS Coverage Database connector

Native real-time access to the Medicare/Medicaid coverage database. Claude pulls Local Coverage Determinations (LCDs) by itself and proposes an authorization outcome *with documented grounds*.

2 • Medical Coding Skill

Maps clinical documentation to ICD-10 and CPT codes, with rationale prepared for coder validation. It doesn't replace the coder — it makes the coder several times faster.

3 • EHR via HealthEx

With patient consent, Claude can query data from Apple Health and Android Health Connect on HIPAA-eligible infrastructure.

4 • Scientific Literature

Native access to PubMed and 10x Genomics datasets. This is a qualitative leap above "paste an abstract and summarize": Claude can **find contradictions between studies**.

5 • Microsoft Foundry

Deployment via Azure for organizations that already have that infrastructure approved by compliance.

90%+ ON MMLU MEDICAL


Both Claude 3.5 Sonnet and Claude 4.5 Opus exceed 90% on medical and biological subsets — a score above many practicing physicians on standardized exams (JAMA Network Open, 04/2026).

But note: the same research reports that **Claude is conservative in escalating diagnostic next steps**. The model is strong at reasoning — cautious at recommending action.

That's a feature, not a bug.

DIAGRAM 02 · ANATOMY OF CLAUDE FOR HEALTHCARE

5 COMPONENTS + COMPLIANCE LAYER



Five components around one core. On the right — a vertical compliance layer (BAA + Zero Data Retention) cutting across every arrow.

— REAL DEPLOYMENTS · NOT DEMOS

Banner Health, Novo Nordisk, Genmab, Novartis. Not pilots — production.

Banner Health — "BannerWise"

33 hospitals, **\$15.6B** in annual revenue, 55,000+ employees. First use case: **oncology chart summarization**. Eight hours of manual physician work per patient → an automated summary for validation.

By end of 2025: **1,400+ clinical notes processed**. 80–85% of users report time savings.

Pharma R&D

- **Novo Nordisk, Genmab** — literature synthesis, hypothesis generation, genomic analysis.
- **Novartis (via Axiom)** — toxicity prediction in preclinical phase. Catching a problem *before* Phase I is orders of magnitude cheaper.
- **Clinical protocols** — e.g., endpoint selection for Parkinson's that's both clinically meaningful and regulator-acceptable.

— MICHAEL REAGIN · CTO, BANNER HEALTH

||

The fact that the model admits when it doesn't know something, instead of inventing a plausible-sounding answer — that's a critical safety property for us in a clinical environment.

BANNER HEALTH · 33 HOSPITALS · \$15.6B IN REVENUE

HIPAA, BAA, Zero Data Retention

— *the most important chapter in this ebook.*

Every deployment that touches PHI has to start **with paperwork, not with code**. It sounds boring. It costs the least. And it's what separates organizations that have an AI deployment from organizations that have a HIPAA incident.

Zero Data Retention — what changed in April 2026

Starting April 2026, the native Anthropic API is classified as an "Eligible Service" **without requiring you to explicitly configure Zero Data Retention**. Practical translation: you can run standard logging and audit alongside Claude *without losing BAA coverage*.

Before, this was a trade-off: turn on logging → lose ZDR. Want auditing → lose BAA. You no longer have to choose.

CLAUDE CODE IS NOW UNDER BAA TOO

Selected Claude Code features are now covered by the BAA. Hospital IT teams can build their own healthcare applications with AI assistance, **without stepping outside compliance**.

⚠ WARNING • CORE INSIGHT

Only two Anthropic products are HIPAA-eligible:

1. Claude.ai Enterprise plan
2. Anthropic API (native, first-party)

Claude Pro and Claude Team are **not** under BAA. Deployments on AWS Bedrock, GCP Vertex, and MS Foundry run under **separate** BAAs with those providers.

If anyone in your organization pastes PHI into Claude Pro — you have a HIPAA incident today.

PRODUCT
BAA
PHI ALLOWED
ZDR
AUDIT LOGS
Claude Pro
✗
✗
✗
✗
Claude Team
✗
✗
✗
✗
Claude.ai Enterprise
✓
✓
✓

✓

Anthropic API (1P)

✓

✓

✓

✓

Deployments via AWS Bedrock / GCP Vertex / MS Foundry operate under those platforms' separate BAAs. Verify scope before deployment.

— CHAPTER 03 • THREE DEPLOYMENT PATTERNS

There is no single
"Claude deployment
in healthcare." There
are *three radically
different patterns*.

CORE INSIGHT

The choice of pattern matters more for project success than the choice of model. Three risk profiles, three cost profiles, three ROI profiles — and three different types of organizations they're the right fit for.

DIAGRAM 04 • THREE PATTERNS SIDE BY SIDE

SPREAD

A

Native Connectors

Claude for Healthcare "out of the box."

YOU
CONNECT

CMS • ICD-10 • HealthEx • PubMed

WHO RUNS IT

Internal team + Anthropic Enterprise

90 DAYS

First production workflow in a narrow use case

RISK

You're constrained to what Anthropic has integrated

COST

\$

B

RCM automation – Flobotics

Claude as the brain inside RPA pipelines.

YOU CONNECT	Epic / Cerner / payer portals · HL7 · FHIR · LIMS
WHO RUNS IT	Automation partner (e.g., Flobotics) + IT
90 DAYS	Hard ROI in weeks — denials, coding, prior auth
RISK	Doesn't change <i>how</i> the org decides — automates existing process
COST	\$\$

C

Data Warehouse

The warehouse is the single source of truth. Zero MCPs.

YOU CONNECT	Read-only DW (Snowflake / BigQuery / Redshift)
WHO RUNS IT	Data + compliance, small team
90 DAYS	Leadership generates its own analytics, T+1
RISK	High Enterprise cost; no real-time
COST	\$\$\$

— PATTERN A

Native Connectors — Claude out-of-the-box.

What it is: you use the connectors Anthropic ships in Claude for Healthcare (CMS, ICD-10, HealthEx, PubMed).

Who it's for: organizations whose workflow fits into "canonical" areas — prior authorization, coding, literature review.

Upsides


- Fastest time-to-value
- Compliance is pre-baked
- Low cost of entry

Downsides

- Limited to what Anthropic integrated
- Doesn't cover your own systems
- Little control over the logic

DIAGRAM 05 · PATTERN A

NATIVE



RCM automation — Claude as the brain inside RPA pipelines.

What it is: Claude is the "brain" inside RPA pipelines. You don't replace Epic/Cerner — **you automate the work *around* them.** FHIR, HL7, payer integration, codes.

What Flobotics ships as the automation partner

- **Prior Authorization** — credential validation, medical-necessity check against coverage policies, traceback to source documentation.
- **Insurance Claims Appeals** — denial review, gap identification, drafted appeal with cited sources.
- **Care Coordination & Triage** — flagging urgent cases, drafting routine responses.
- **Ambient Scribing** — visit recording → complete documentation with coding in ~47 seconds.
- **Medical Coding** — code validation against documentation and policies *before* a claim is submitted.

Upsides

No replatforming. ROI in weeks, not quarters.

Downsides

Doesn't fundamentally change *how* the org decides — it automates existing process.

SOURCE SYSTEMS

Epic

Cerner

HL7 / FHIR

Payer portals

RPA + CLAUDE · ORCHESTRATOR

Validate · Decide · Draft

OUTPUTS

Prior auth
decisions

Appeal drafts
with citations

Coded
visits (ICD/CPT)

NUMBERS FROM REAL FLOBOTICS DEPLOYMENTS

~90%

reduction in manual rework in denials
management

6 min → 30

S

processing time per record

80%

of a 3,000-record backlog cleared in
3 weeks

\$90K+

savings, 500+ hrs/month returned to
the org

70%

faster claim turnaround time

98%

reduction in manual claim prep (~250
hrs/wk)

Sources: Flobotics case studies — Gentem (98% reduction), PathGroup (denials throughput), PTCOA (ROI in 23 days), Pain Treatment Centers of America.

The Data Warehouse as the single source of truth.

Zero MCPs to source systems.

This is the pattern nobody shows on webinars. And the one a large diagnostic lab just deployed.

What it is: you don't connect Claude to source systems *at all*. No MCP to Xifin. No MCP to Salesforce. No MCP to Careviso. No MCP to the LIS.

Instead:

1. All source systems load into the warehouse (once a day).
2. Claude Enterprise gets **read-only** access to the warehouse.
3. Leadership talks to Claude in natural language and generates analytics themselves.

PROJECT PREMISE

"Let leadership build their own analytics with Claude."

DIAGRAM 07 · PATTERN C — FLAGSHIP

ZERO MCP · READ-ONLY · DAILY REFRESH

SOURCE SYSTEMS



ETL · daily load



Data Warehouse
Snowflake / BigQuery / Redshift

read-only



C
Claude Enterprise
BAA · Audit · ZDR

Zero MCP to source systems · Read-only contract · Daily refresh

Leadership · self-service

WHY THIS WORKS EVEN THOUGH IT LOOKS LIKE "A STEP BACK"

- **Attack surface = one database.** You audit one integration, not ten.

- **Stable schema.** Claude doesn't guess what "customer" means — it knows, because the warehouse has canonical dimensions.
- **Once-a-day freshness is enough.** Leadership needs consistency, not real-time.
- **No side effects.** Read-only = Claude can't break Xifin or Salesforce — not even hypothetically.

WHEN THIS IS THE RIGHT PATTERN

- You have many source systems written across different decades.
- You already have a data warehouse (BigQuery / Snowflake / Redshift).
- Your user persona is a **decision-maker**, not an operator.
- The org tolerates "yesterday at 11:59 PM" instead of "now".
- You treat compliance as architecture, not a checklist.

Cost: Claude Enterprise is expensive. Very expensive. In return you get something you cannot buy in Claude Team: **BAA, audit trail, and one point of legal accountability.**

— CHAPTER 04 • CASE STUDY

A large diagnostic lab. Four systems.

Zero MCPs.

Context

Four primary systems: **Xifin** (lab RCM), **Salesforce** (CRM/sales), **Careviso** (prior auth / benefits investigation), and a homegrown **Lab Management System**.

The problem: leadership asks questions nobody can answer in under 3 days — because the answer needs data stitched across four systems, and each system has its *own* definition of "customer," "visit," and "revenue."

Architectural decision

Deliberate rejection of two "trendy" options:

- ✗ **No:** a separate MCP/connector for every source system.
- ✗ **No:** real-time streaming of data into Claude.
- ✓ **Yes:** the warehouse as a single contract, Claude Enterprise as the sole AI access point.

WHAT THEY DELIBERATELY SACRIFICED

- Data freshness (T+1 instead of real-time)
- Claude's ability to "write" back to source systems
- Lower cost (Enterprise is expensive)

WHAT THEY DELIBERATELY GAINED

- One BAA, one audit log
- Stable schema → fewer hallucinations
- Self-service analytics for non-technical leadership
- No "careful, AI is writing to Salesforce" pressure

The VP of Sales asks a question. Response time: *minutes, not days.*

● Claude Enterprise · BAA · Read-only to DW

Diagram 08

Show me patients with claims denied in Q1 2026 where Careviso had an active benefit but Xifin failed to capture the code.

Found 217 such records (Q1 2026). Three main patterns:

```
patients_with_active_benefit_but_denied_claim
├─ 117  · missing CPT code on claim line  (Xifin)
├─ 68   · POS code mismatch              (Xifin ↔ Careviso)
└─ 32   · benefit active after service date (Careviso)
```

Anomaly: 117 of 217 (54%) are the same payer — UnitedHealthcare. In January they changed the POS requirement; the ETL integration likely missed the policy update.

3 hypotheses:

1. Process gap in Xifin: no POS validation before claim write.
2. ETL bug (payer policy change on Jan 14).
3. Outdated code in payer policy — needs escalation to UHC.

Want me to pull the 117 records as a CSV? Or a breakdown by provider?

— LEAD ARCHITECT · DIAGNOSTIC LAB

||


The most expensive AI deployment we've done. And the only one that cleared the compliance audit without a single finding.

CASE STUDY · CLAUDE ENTERPRISE + DATA WAREHOUSE · 2026

— CHAPTER 05 • ROI AND PATTERN SELECTION

The decision matrix. Two axes, four answers.


— DIAGRAM 09 • PATTERN SELECTION MATRIX



The selection matrix is mundane, but it works because it forces you to name two things: **how complex your system landscape is** and **how fast you have to show a return**.

Most organizations make the same mistake: they start from the model (“we want to deploy Claude”) instead of from the bottleneck (“we have 3,000 denials in queue and 6 minutes per record”). The matrix flips that order.

Three areas where Pattern B delivers documented ROI in under 90 days:



Percentages = reduction in time / manual work / backlog, based on Flobotics deployments across 50+ U.S. healthcare providers.

RCM · NEAR-TERM ROI

Organizations deploying automation comparable to Flobotics report 15–30% reduction in denial rate in 12–18 months. Prior authorization and medical coding have documented returns.

LIFE SCIENCES R&D · LONGER HORIZON

Genomic synthesis and preclinical work have higher magnitude potential, but deployment complexity and regulatory validation time aren't comparable to back-office automation.

— CHAPTER 06 • SIX PITFALLS

Six decisions that cost twice.

HEADS-UP

Each of the below looks like a shortcut. Each one leaves the organization with a deployment that's more expensive, slower, and less secure than it could have been.

01 "Let's give people Claude Pro and let them try it."

The first pasted PHI = a HIPAA incident. **Only Enterprise and the native API are under BAA.** Everything else is legal risk you won't see until an audit comes.

02 "We'll wire up every source system via MCP."

You're multiplying your attack surface and integration failure points. Often the warehouse is the right answer. Try starting with one contract, not ten.

03 "Model first, use case later."

Wrong way around. Start from the bottleneck (denials? coding? leadership analytics?), **then pick the pattern, then pick the model.** Otherwise you're buying a solution looking for a problem.

04 "Real-time or nothing."

Leadership almost never needs real-time. They need **consistency**. T+1 with stable schema beats T+0 with ten inconsistent sources.

05 "Claude will handle uncertainty on its own."

It will only if the prompts and system encourage it to say "I don't know." The most common cause of hallucinations is a system prompt that rewards decisive answers.

"BAA = compliance done."

BAA is a foundation, not a ceiling. You still need: audit log, retention, role-based access, training, incident response.

Miss any of those and you don't have compliance — you have paperwork.

— CHAPTER 07 • THE FIRST 90 DAYS

Three phases. Three decisions. One path to production.

DAYS 0–30 • FOUNDATION

- ☐ Decision: Enterprise plan or native API
- ☐ BAA signed
- ☐ Map of source systems touching PHI
- ☐ Pattern selection (A / B / C)
- ☐ Owner on the business side and on IT

DAYS 31–60 • FIRST USE CASE

- ☐ One narrowly scoped bottleneck (e.g., *Q1 denials, payer X*)
- ☐ Baseline: time, cost, error rate
- ☐ Pilot with one team
- ☐ Prompts + evaluation + guardrails
- ☐ Audit log on from day zero

DAYS 61–90 • MEASURE & DECIDE

- ☐ Compare to baseline (hard numbers)
- ☐ Audit log review
- ☐ Honest: where does Claude get it wrong?
- ☐ Scaling plan *or* conscious “no-go”



Sponsor for the second use case

90 days isn't a deadline for "a working product for the whole org." It's a deadline for **a hard decision**: do we scale this pattern, or do we deliberately stop? Both decisions are legitimate — indecision is not.

— CONCLUSION

2026 isn't the question of "AI in healthcare." It's the question of *which of the three patterns*.

Native Connectors are fastest but limited. RPA + Claude delivers hard ROI in weeks. The Data Warehouse is the most expensive and most secure — and that's exactly why it's chosen by organizations that understand **compliance architecture isn't a checklist, it's a foundation**.

The common denominator across all three: **Anthropic Enterprise or the native API, BAA before the first prompt, audit log from day zero**.

The rest — pattern selection, use case sequencing, the cost curve — depends on your landscape. And that's exactly the conversation that doesn't go well over a deck.

CONSULTATION WITH AN AI DEPLOYMENT ADVISOR

Want to validate which pattern fits your organization?

Karl Mielnicki, Co-founder & CTO at Flobotics, runs 45-minute consultation sessions. The output: a written pattern recommendation, a list of three bottlenecks with the highest ROI, and a deployment cost estimate.

Book 45 minutes →

k.mielnicki@flobotics.io · +1 (910) 518-0124

ABOUT THE AUTHOR

Karl Mielnicki

Co-founder & CTO, Flobotics.



Karl has been building AI agents and custom automation for healthcare RCM since 2019. He has worked with 50+ U.S. providers and payers — from single clinics to hospital networks. Previously in process automation and medical-systems integration.

Writes and speaks on ROI in healthcare AI — featured by **Forbes** (3×).

Book a call →

flobotics.io

ABOUT FLOBOTICS

Flobotics builds **AI agents** and custom automation for healthcare RCM. End-to-end: from eligibility checks, through prior auth, to claims and denials management. **100% HIPAA compliant**. No API required.

Reach us

k.mielnicki@flobotics.io

+1 (910) 518-0124

Made with love in Indiana, US.

More resources

- Case studies
- ROI calculator
- Blog
- RCM Glossary

SOURCES

1. *Claude for Healthcare Overview · 2026* — flobotics.io/blog/claude-for-healthcare
2. *Claude for Healthcare Automation* — flobotics.io/claude-for-healthcare-automation
3. JAMA Network Open, April 2026 — MMLU Medical benchmarks
4. J.P. Morgan Healthcare Conference, January 2026 — Anthropic announcement
5. Internal case study, diagnostic lab (Claude Enterprise + Data Warehouse), 2026
6. Flobotics — case studies: Gentem (98% reduction in claim prep), PathGroup (denials), Pain Treatment Centers of America (ROI in 23 days), SuperBill (CTT 70%)

Disclaimer: This ebook is not legal advice on HIPAA, GDPR, or any other regulation. Consult your organization's compliance officer before deployment. Numbers from Flobotics deployments reflect specific projects and may not be representative of every organization.



© 2026 Flobotics. Made with love in Indiana, US. · Designed for healthcare leaders, not for HIPAA officers (but they will be happy too).