ENTERPRISE PRODUCT REQUIREMENTS DOCUMENT

HelixCare AI Billing Assistant

A portfolio-ready enterprise PRD for a fictional healthcare network deploying an LLM-powered conversational AI product for billing inquiries, balance explanations, payment routing, escalation, and post-launch optimization.

PRODUCT NAME
HelixCare AI Billing Assistant

PREPARED BY
Simone Wilson

VERSION
v1.0 · Portfolio Edition

SANITIZED PORTFOLIO SAMPLE · FICTIONAL CLIENT

01

Executive Summary

HelixCare AI Billing Assistant is an enterprise conversational AI product designed to help patients understand balances, verify account information, review payment options, resolve common billing questions, and escalate complex financial concerns to the right support team.

Product Thesis

→ The product is not a novelty chatbot. It is a workflow orchestration layer for high-volume, sensitive, and operationally expensive healthcare billing interactions.

Business Value

→ Reduce repetitive billing support volume.

→ Improve patient clarity around balances and payment options.

→ Increase escalation accuracy for sensitive billing scenarios.

→ Standardize billing interactions across locations and channels.

PRODUCT POSITIONING

Create a trusted AI billing experience that helps patients resolve routine billing needs quickly while routing complex or sensitive issues to the right human support path.

02

Business Problem

Healthcare billing interactions are often confusing, emotionally charged, and operationally expensive. The product opportunity lives at the intersection of patient clarity, support efficiency, and compliant workflow design.

Patient Problem

- → Patients struggle to understand balances and billing language.
- → Patients do not always know whether insurance is pending or payment is required.
- → Patients need reassurance, plain language, and a fast human path when issues are sensitive.

Operational Problem

- → Billing teams handle high volumes of repetitive balance and payment questions.
- → Manual lookups increase handle time and reduce throughput.
- → Escalation pathways are inconsistent across locations and issue types.

Business Problem

- → Avoidable support costs rise as billing volume increases.
- → Payment resolution may be delayed by confusion or routing friction.
- → Inconsistent billing experiences can reduce trust and satisfaction.

03

Product Vision, Goals & Non-Goals

The product must improve billing self-service without sacrificing trust, compliance, or the human support path required for sensitive financial scenarios.

PRODUCT VISION

Create a trusted AI billing experience that helps patients resolve routine billing needs quickly while routing complex or sensitive issues to the right human support path.

Goals

- → Reduce repetitive billing support volume.
- → Improve patient self-service completion.
- → Increase clarity around balances and payment options.
- → Improve escalation accuracy.
- → Standardize billing workflows across locations.
- → Provide secure, compliant, patient-friendly interactions.

Non-Goals

- → Replace billing specialists.
- → Make financial hardship determinations without human review.
- → Provide legal or insurance coverage advice.
- → Expose internal minimum-payment rules.
- → Resolve disputed charges without escalation.
- → Override existing compliance or payment policies.

04
Users & Stakeholders

The product must serve patients while also reducing operational burden for billing teams, revenue cycle leaders, and engineering/integration teams.

Primary Persona: Patient With Billing Question

→ Needs clarity, reassurance, fast routing, payment options, and confidence that sensitive information is protected.

Secondary Persona: Billing Support Agent

→ Needs fewer repetitive calls, cleaner handoffs, interaction context, and clearer escalation reasons.

Internal Persona: Revenue Cycle / Operations Leader

→ Needs efficiency, compliance, consistent workflows, measurable outcomes, and visibility into unresolved issues.

Technical Persona: Engineering / Integration Team

→ Needs clear API requirements, data dependencies, error states, logging rules, and acceptance criteria.

05

Current-State & Future-State Workflow

The workflow redesign shifts the experience from queue-heavy manual handling to secure AI-guided orchestration with clear escalation and auditability.

Current-State Workflow

01	**Patient receives bill**	Confusion or concern prompts a call.
02	**Patient waits in queue**	Billing intent is not always captured early.
03	**Agent verifies identity**	Manual lookup required before account details are shared.
04	**Agent explains options**	Explanation varies by agent, policy, and account context.
05	**Payment or escalation**	Transfer path may be inconsistent or unresolved.

Future-State AI Workflow

01	**AI identifies intent**	Balance, payment, dispute, hardship, or live agent request.
02	**Patient verification**	Secure identity check before account-specific disclosure.
03	**Balance context retrieval**	Current balance, due date, payment status, and insurance-pending flag where available.
04	**Approved explanation**	Patient-friendly summary using retrieved data only.
05	**Payment or handoff**	Routes to payment system or human agent with reason and context.

06

Core Product Capabilities

Capabilities are modular so the product can launch with a focused billing use case and expand into additional revenue cycle workflows over time.

Intent Recognition Recognizes balance inquiry, make payment, payment plan question, insurance confusion, dispute, financial hardship, and live agent requests.	**Patient Verification** Verifies identity using approved identifiers such as phone number, DOB, account number, ZIP code, and last name before disclosure.
Balance Explanation Explains current balance, due date, statement date, payment status, and insurance-pending status using plain, approved language.	**Payment Routing** Supports card payment, bank account payment, payment portal routing, and agent-assisted payment paths when configured.
Escalation Management Escalates disputed charges, hardship concerns, failed verification, frustrated callers, low confidence responses, and compliance-sensitive scenarios.	**Audit & Reporting** Captures intent, outcome, escalation reason, unresolved issue, completion status, failure point, and channel metadata.

07

Functional & Non-Functional Requirements

Requirements are written to support engineering handoff, QA/UAT planning, compliance review, and post-launch performance evaluation.

Functional Requirements

ID	REQUIREMENT	ACCEPTANCE CRITERIA
FR-001	Intent Capture	The system must identify the patient's primary billing intent within the first two interaction turns.
FR-002	Identity Verification	The system must verify patient identity before disclosing account-specific billing information.
FR-003	Balance Retrieval	The system must retrieve current balance data from the billing system before presenting account-specific information.
FR-004	Plain-Language Explanation	The system must explain billing information using patient-friendly, non-technical language.
FR-005	Payment Option Presentation	The system must present only approved payment options configured by the organization.
FR-006	Escalation Triggering	The system must escalate to a human agent when confidence falls below threshold or when the intent requires human review.
FR-007	Interaction Logging	The system must log interaction outcome, escalation reason, and completion status.

Non-Functional Requirements

Security & Compliance

- → HIPAA-aware handling of PHI/PII.
- → No unnecessary exposure of account-specific information.
- → Approved language only for billing and payment explanations.
- → Escalation for legal, insurance, hardship, or dispute scenarios.

Reliability & Performance

- → Graceful fallback when APIs fail or data is unavailable.
- → No hallucinated account information.
- → Clear fallback when confidence is low.
- → Voice-friendly prompts and plain-language responses.

08

AI Behavior & Guardrails

This is the AI-specific layer of the PRD. It defines how the system behaves when accuracy, trust, policy, or patient safety is at stake.

AI Must

- → Use approved billing language.
- → Verify identity before account-specific disclosure.
- → Summarize only retrieved source-system data.
- → Escalate uncertainty or sensitive issues.
- → Acknowledge limitations when data is unavailable.
- → Maintain a calm, patient-friendly tone.

AI Must Not

- → Invent balances, due dates, payment status, or account details.
- → Disclose account information before verification.
- → Negotiate unauthorized payment arrangements.
- → Give legal or insurance coverage determinations.
- → Override policy or compliance instructions.
- → Pressure patients or expose internal rules.

09
API & Data Dependencies

The PRD intentionally separates known product needs from open technical questions, reducing the risk of inventing API behavior before discovery confirms the source of truth.

SYSTEM	REQUIRED DATA / FUNCTION	PRODUCT USE
Patient Account System	Patient identity, account status, contact information	Verification, account matching, handoff context
Billing System	Current balance, due date, statement date, payment status, insurance pending flag	Balance explanation and payment eligibility
Payment Processor	Payment link, available methods, transaction confirmation	Payment routing and confirmation
CRM / Ticketing	Escalation ticket, interaction log, agent handoff context	Human handoff and operational tracking
Telephony / Chat	ANI, session metadata, transcript, channel source	Personalization, logging, QA, and analytics

OPEN TECHNICAL QUESTIONS

Which source of truth owns balance data? Are payment plan rules exposed via API? What error codes are returned when balance data is unavailable? Which fields can be safely displayed in the AI channel?

10

Edge Cases & Exception Handling

Exception handling is a senior product signal. The product must know what to do when the happy path breaks, because in healthcare billing, the happy path is only one creature in the forest.

EDGE CASE	SYSTEM BEHAVIOR	ESCALATION / LOGGING
Patient fails verification	Do not disclose account-specific information. Offer alternate verification or agent handoff.	Log failed verification reason and handoff path.
API timeout	Apologize, explain account details are temporarily unavailable, and offer agent transfer.	Log API failure and affected system.
Multiple accounts found	Ask approved clarifying question without exposing details.	Escalate if match remains ambiguous.
Account sent to collections	Use approved language and route to appropriate support path.	Escalate to billing specialist or collections team.
Patient disputes bill	Acknowledge concern and avoid making determinations.	Escalate with dispute reason and transcript context.
Patient expresses hardship	Use empathetic language and avoid unauthorized promises.	Escalate to approved financial assistance workflow.
Caller is not the patient	Follow authorized representative workflow if available.	Escalate if authorization cannot be confirmed.
AI confidence is low	Stop account-specific explanation and route safely.	Log low-confidence intent and escalation reason.

11

Success Metrics

Metrics are grouped by product performance, operational efficiency, patient experience, and AI quality so leadership can evaluate impact beyond launch completion.

Product	Ops	Experience	AI Quality
Task completion rate, containment rate, successful payment routing, first-contact resolution.	Billing call volume reduction, average handle time reduction, escalation accuracy, agent workload reduction.	Patient satisfaction, repeat contact rate, abandonment rate, complaint rate.	Intent accuracy, hallucination incidents, fallback rate, low-confidence escalation rate.

12

Launch Plan, Risks & Open Questions

A phased launch limits risk, validates workflow assumptions, and creates the operational feedback loop required for scale.

01 — Discovery & Workflow Validation
Document billing workflows, approved language, escalation paths, API dependencies, and compliance requirements.

02 — Prototype
Build IVA/chat flow for balance inquiry, identity verification, and payment routing.

03 — Pilot
Deploy to a limited patient segment, location group, or low-risk billing workflow.

04 — QA / UAT
Test edge cases, compliance language, API failures, routing paths, and handoff quality.

05 — Rollout
Expand by location, workflow type, patient population, or channel after pilot stabilization.

06 — Optimization
Tune prompts, intents, escalation logic, reporting, and knowledge base based on real interaction data.

Risks & Mitigations

RISK	IMPACT	MITIGATION

Incorrect balance disclosure	Compliance risk	Verify identity and retrieve only source-system data.
API unavailable	Poor experience	Graceful fallback and agent escalation.
Patient distrusts AI	Low adoption	Use clear language and offer easy human handoff.
Poor escalation design	Frustration	Define escalation rules by intent and confidence.
Unauthorized payment guidance	Compliance risk	Use approved payment policy language only.

Open Questions

Product & Strategy

- → What is the source of truth for current balance?
- → Which billing scenarios require human review?
- → Are payment plans supported through API?
- → Can payment links be generated dynamically?

Compliance & Operations

- → What verification fields are approved?
- → Which languages are required at launch?
- → What reporting fields are needed by operations?
- → What compliance language must be approved?

A

Appendix

A compact proof layer showing conversation design, user story thinking, and escalation structure.

Sample Conversation

- → AI: Thanks for calling HelixCare Billing. I can help with balance questions, payment options, or connect you with a billing specialist. What can I help you with today?
- → Patient: I got a bill and I don't understand what I owe.
- → AI: I can help explain your balance. Before I share account details, I'll need to verify your identity.

Sample User Stories

- → As a patient, I want to understand my current balance so I can decide what to do next.
- → As a patient, I want to know my payment options so I can choose the path that works for me.
- → As a billing agent, I want handoff context so I do not ask the patient to repeat information.
- → As an operations leader, I want outcome reporting so I can measure self-service and escalation performance.
- → As a compliance stakeholder, I want guardrails so the AI does not disclose information before verification.

Prepared by Simone Wilson
Senior AI Product Manager · Enterprise AI Transformation & Conversational AI Systems
simonewrites40@gmail.com · linkedin.com/in/yunchewilson