
Riverside Family Medicine

As of September 21, 2026

BUSINESS PROFILE

Multi-site family medicine group serving 18,000 active patients across four clinics. Services include primary care, chronic disease management, and same-day sick visits. Front-desk and nursing staff coordinate referrals, prior authorizations, and follow-up outreach.

BUSINESS TYPE	INDUSTRY
Private Company	Healthcare – Providers & Services
LOCATION	OPERATING REGIONS
United States	North America
COMPANY SIZE	CUSTOMER TYPES
Medium (51–250 employees)	Consumer
BUSINESS MODEL	
Fee-for-Service	
REGULATORY CONTEXT	
HIPAA, Other	
STRATEGIC PRIORITIES	
Operational Efficiency, Compliance, Talent, Cost Reduction	
KEY SYSTEMS	
Custom / Proprietary, HRIS, Accounting	
DATA SENSITIVITY	
Protected Health Information (PHI), Personally Identifiable Information (PII)	

CURRENT BUSINESS AI POSTURE

MARKETING	SALES
Experimenting	Avoiding
Marketing coordinator uses ChatGPT for patient education blog drafts. No clinical review workflow.	New patient intake is fully manual; no AI in scheduling or eligibility checks.
DELIVERY	CUSTOMER SUCCESS
Experimenting	Avoiding
Two physicians trial ambient scribe tools in pilot rooms. Not approved for full clinic rollout.	Care-gap outreach and recall letters are template-based; no AI drafting.
OPERATIONS	GROWTH
Experimenting	Avoiding
Admin staff use ChatGPT for non-clinical email templates. PHI is explicitly prohibited by policy but not audited.	Quality improvement and credentialing workflows have no AI tooling.
IT / TECHNOLOGY	LOGISTICS

Avoiding

EHR and practice-management systems are vendor-managed. No AI tooling in IT or security operations.

Avoiding

Medical supply ordering is managed by the office manager via distributor portals. No AI-assisted procurement.

CURRENT AI USE

Patient education blog drafts	
BUSINESS FUNCTION	TOOL / VENDOR
Marketing	ChatGPT
WHO IS USING	WORK SUPPORTED
Marketing coordinator	Drafting seasonal wellness articles for the clinic website
PROCESS DOCUMENTED	IS IT HELPING?
Informal / primarily in one person's head	Yes
CREATING REWORK / PROBLEMS?	
No	
INFORMATION SHARED	
No PHI — generic health education topics only.	
APPROVED?	
Yes	
NOTES	
Medical director reviews before publish.	

Ambient visit documentation pilot	
BUSINESS FUNCTION	TOOL / VENDOR
Delivery	Pilot scribe vendor (evaluation)
WHO IS USING	WORK SUPPORTED
2 physicians in pilot exam rooms	Drafting visit summaries from clinician-patient conversation
PROCESS DOCUMENTED	IS IT HELPING?
Documented and repeatable	Partly
CREATING REWORK / PROBLEMS?	
Partly	
INFORMATION SHARED	
PHI processed under BAA with pilot vendor only.	
APPROVED?	
Yes	

NOTES

Compliance review required before expanding beyond pilot rooms.

Non-clinical staff email templates

BUSINESS FUNCTION

Operations

TOOL / VENDOR

ChatGPT

WHO IS USING

Front-desk supervisors

WORK SUPPORTED

Scheduling confirmation and general office correspondence

PROCESS DOCUMENTED

Informal / primarily in one person's head

IS IT HELPING?

Yes

CREATING REWORK / PROBLEMS?

No

INFORMATION SHARED

No PHI permitted — generic templates only.

APPROVED?

Partly

NOTES

Policy exists but adherence is not routinely audited.

Scheduling No-Show Prediction

Operations · In Progress

OPPORTUNITY BEING ASSESSED

TITLE

Scheduling No-Show Prediction

BUSINESS FUNCTION

Operations

OWNER

James Okonkwo, Clinical Operations Manager

STATUS

Deferred

VALUE PROPOSITION

Improve resource utilization

CURRENT FRICTION

Same-day cancellations leave open slots. Front desk overbooks reactively, creating long wait times on heavy days.

DESIRED CHANGE

Use de-identified operational signals to flag no-show risk for human scheduling review without automating patient access decisions.

SUCCESS INDICATOR

Pilot report identifies high-risk slots with measurable precision; wait times do not increase; no demographic group receives worse reminder coverage.

AI'S ROLE

AI produces no-show risk reports and suggested reminder queues. Front-desk supervisors decide any scheduling changes.

ASSESSMENT

STATUS

In Progress

VERSION

LIGHTNING_1.0

STARTED

9/21/2026

READINESS PROFILE

Dimension	Status	Key Observation
Process	NEEDS ATTENTION	Scheduling rules vary by provider, visit type, and season. Before AI predicts no-shows or suggests overbooking, the clinic needs a documented decision path for who approves schedule changes.
Data & Information	NEEDS ATTENTION	Appointment history is available, but no-show reasons, patient communication preferences, and provider constraints are inconsistently captured. The model should start with de-identified operational fields, not clinical notes.
People & Training	NEEDS ATTENTION	Front-desk staff will need guidance on using AI suggestions as prompts, not mandates. Operations leadership should define when human judgment overrides a prediction.
Risk & Security	NEEDS ATTENTION	Overbooking based on weak predictions could increase wait times and patient frustration. Equity review is needed so reminders or slot changes do not disadvantage specific patient groups.
Technology & Vendor	NEEDS ATTENTION	Scheduling data lives in the EHR/practice-management system, but there is no tested prediction workflow. A reporting-only pilot should precede any automated schedule recommendations.

OVERALL INTERPRETATION

Scheduling optimization should wait until clinical documentation pilots stabilize. Front-desk AI use must remain strictly non-clinical until policy training is complete.

QUESTION RESPONSES

PROCESS

We can explain how the current work is performed from beginning to end.	YES
The people involved generally perform the work in a consistent way.	YES
We understand what a good result looks like.	YES
We know where judgment, exceptions, or escalation are required.	YES
The process is documented well enough to support testing or improvement.	YES
We know who owns the process and who can approve changes to it.	YES

DATA & INFORMATION

We know what information the AI would need to perform this work.	PARTLY
We know where that information comes from and who owns it.	PARTLY
The information is accurate, complete, and usable enough for the intended purpose.	NO
We know what information would be shared with the AI and understand whether it is appropriate and permitted to share that information with the specific tool being considered.	NO
We understand whether customer, employee, financial, confidential, proprietary, regulated, or restricted information is involved.	PARTLY
We know what information should not be provided to the AI.	PARTLY

PEOPLE & TRAINING

The people involved are willing to try AI for this use case.	YES
They can learn how to provide useful instructions and context.	YES
The proposed AI use fits reasonably into how the work is performed.	YES
Someone involved has enough subject-matter expertise to evaluate the AI output.	PARTLY
We know what quality or accuracy is acceptable for this use case.	PARTLY
People understand that AI output may need to be challenged, corrected, or rejected.	PARTLY
It is clear who remains responsible for the final work or decision.	PARTLY
People understand when human review is required.	PARTLY

RISK & SECURITY

We know whether this use is internal, employee-facing, customer-facing, vendor-facing, or public.	YES
We know who will see, use, or rely on the AI output.	YES
We understand what could happen if the AI is wrong.	YES
We know when human review is required before the output is used.	PARTLY
We have or can establish acceptable-use boundaries for this use case.	PARTLY
We can document, when appropriate, what the AI did, what information it used, who reviewed it, and what action was taken.	PARTLY

We know who is responsible for identifying and responding to security or privacy concerns.	PARTLY
--	--------

TECHNOLOGY & VENDOR

We know which technology platforms and vendors touch this process or its information.	PARTLY
We know whether those vendors use AI in relevant features or services.	PARTLY
We understand what information may be sent to or processed by the AI service.	NO
We understand how information is stored, retained, or shared.	NO
We know whether information may be used to train or improve models or services.	NO
We know what controls are available, including whether AI features can be disabled or configured.	NO
We understand who has access to the data and what happens to it if we stop using the vendor.	NO
The proposed technology is a reasonable fit for the intended business use.	NO

OPEN QUESTIONS

Which non-clinical scheduling fields can be used for prediction without relying on clinical notes?	MEDIUM	OPEN
Who approves an overbooking or reminder change suggested by the model?	MEDIUM	OPEN
How will the clinic monitor whether predictions disadvantage specific patient groups?	MEDIUM	OPEN

ASSUMPTIONS

Initial pilot is reporting-only and does not automatically change schedules.	UNVERIFIED
Front-desk supervisors approve any slot or reminder action.	UNVERIFIED
Prediction inputs exclude clinical note content.	UNVERIFIED

RISKS

Weak predictions cause overbooking and longer patient wait times	OPEN
LIKELIHOOD	DECISION
Medium	Reduce
Reminder targeting disadvantages a patient group or communication preference	OPEN
LIKELIHOOD	DECISION
Medium	Reduce
Staff treat no-show scores as mandates instead of prompts	OPEN

LIKELIHOOD	DECISION
Medium	Reduce

CONSTRAINTS

Scheduling workflows must avoid discriminatory patient-access impacts.	REGULATORY	IDENTIFIED
No-show reasons and communication preferences are not consistently captured today.	OTHER	IDENTIFIED

ACTION ITEMS

Document scheduling decision rules and override authority	GOVERNANCE	OPEN
Create de-identified scheduling export for reporting-only pilot	PILOT PREPARATION	OPEN
Define equity review metrics for reminders and slot changes	GOVERNANCE	OPEN
Pilot weekly no-show risk report before any overbooking changes	PILOT PREPARATION	OPEN

BUSINESS OBJECTIVES

Identify high-risk appointment slots with useful precision before making schedule changes
EXPECTED VALUE
Other
Maintain average patient wait time within baseline range during pilot
EXPECTED VALUE
Quality Improvement

TEAM READINESS

Name	Role	AI Familiarity	Business Knowledge	Process Knowledge	Review Capability	Risk Awareness	Notes
James Okonkwo	Clinical Operations Manager	—	—	—	—	—	—
Maya Lewis	Front Desk Supervisor	—	—	—	—	—	—

Name	Role	AI Familiarity	Business Knowledge	Process Knowledge	Review Capability	Risk Awareness	Notes
Sarah Chen	Compliance & Privacy Officer	—	—	—	—	—	—

AI ROLE ASSIGNMENTS

Role	Person	Notes
AI Sponsor	James Okonkwo	—
AI Champion	Maya Lewis	—
AI Reviewer	Sarah Chen	—

OPERATING GUARDRAILS

No automated schedule changes during the first pilot phase.	HUMAN REVIEW	PROPOSED
Prediction inputs must exclude clinical notes and use only approved operational fields.	DATA HANDLING	PROPOSED
Equity review is required before any reminder or overbooking workflow expands.	OTHER	PROPOSED

PILOT CHARTERS

Same-day cancellations and reactive overbooking create idle provider time on some days and long waits on others.	DRAFT
USERS	
Front-desk supervisor and scheduling team for four clinics	
IN SCOPE	
Weekly no-show risk report, reminder queue suggestions, human-approved scheduling experiments	
OUT OF SCOPE	

Automated overbooking, clinical triage, patient eligibility decisions

INFORMATION

Appointment history, visit type, provider schedule, prior no-show flags, communication preferences

WORKFLOW

1. Export approved scheduling fields. 2. AI produces no-show risk report. 3. Supervisor reviews suggested reminders. 4. Human-approved actions are logged. 5. Wait-time and equity metrics are reviewed.

HUMAN REVIEW APPROACH

Front-desk supervisor approval required for every scheduling or reminder action.

SUCCESS MEASURES

Useful high-risk slot identification; wait times within baseline; no adverse equity signal

RISKS SUMMARY

Overbooking errors, patient frustration, inequitable reminder targeting

CONTROLS

Reporting-only phase; supervisor approval; equity metrics; no clinical-note inputs

OWNER

James Okonkwo

DURATION

60 days

STOP CONDITIONS

Wait times increase materially; equity review flags harm; staff bypass supervisor approval

DECISION POINT

Day 30: continue reporting-only or stop. Day 60: test limited human-approved scheduling changes or defer.

Referral Letter & Prior Auth Drafts

Delivery · In Progress

OPPORTUNITY BEING ASSESSED

TITLE

Referral Letter & Prior Auth Drafts

BUSINESS FUNCTION

Delivery

OWNER

James Okonkwo, Clinical Operations Manager

STATUS

Assessing

VALUE PROPOSITION

Reduce errors and rework

CURRENT FRICTION

Specialist referral letters and prior authorization narratives are retyped from scratch. Nursing staff chase missing clinical details.

DESIRED CHANGE

Create first drafts from signed visit data and payer-required fields so nurses spend less time retyping and physicians review clinical rationale before submission.

SUCCESS INDICATOR

Referral prep time reduced by 30%; denial-related rework does not increase; 100% physician review on medical-necessity language.

AI'S ROLE

AI assembles referral and prior-authorization drafts from approved EHR fields and payer templates. Nurses review completeness, and physicians approve clinical rationale.

ASSESSMENT

STATUS

In Progress

VERSION

LIGHTNING_1.0

STARTED

9/21/2026

READINESS PROFILE

Dimension	Status	Key Observation
Process	READY TO INVESTIGATE	Referral and prior authorization prep has a known path: signed visit data, required payer fields, nursing review, and physician approval when clinical judgment is involved. AI can help assemble first drafts, but it should not decide medical necessity.
Data & Information	SIGNIFICANT CONCERN	Referral packets and prior-auth narratives contain PHI, diagnoses, medications, and payer-specific details. Any pilot needs BAA coverage, role-based access, and a minimum-necessary data template.
People & Training	NEEDS ATTENTION	Nursing staff know where details are missing, but they need clear rules for when to return a draft to the physician versus completing administrative fields themselves.
Risk & Security	SIGNIFICANT CONCERN	Incorrect clinical rationale could delay care or create denial risk. Physician review is required for medical-necessity language, and compliance should audit early samples.
Technology & Vendor	NEEDS ATTENTION	EHR and payer portals are not connected to a governed drafting workflow. Begin with structured EHR-field extracts and approved document templates before pursuing portal automation.

OVERALL INTERPRETATION

Referral letter drafting is a strong secondary use case once visit summary accuracy is proven. Nursing staff would benefit most from structured drafts that pull from signed visit data rather than retyping narratives.

QUESTION RESPONSES

PROCESS

We can explain how the current work is performed from beginning to end.

YES

The people involved generally perform the work in a consistent way.

YES

We understand what a good result looks like.

YES

We know where judgment, exceptions, or escalation are required.

YES

The process is documented well enough to support testing or improvement.

YES

We know who owns the process and who can approve changes to it.

YES

DATA & INFORMATION

We know what information the AI would need to perform this work.

PARTLY

We know where that information comes from and who owns it.

PARTLY

The information is accurate, complete, and usable enough for the intended purpose.

NO

We know what information would be shared with the AI and understand whether it is appropriate and permitted to share that information with the specific tool being considered.

NO

We understand whether customer, employee, financial, confidential, proprietary, regulated, or restricted information is involved.

PARTLY

We know what information should not be provided to the AI.

PARTLY

PEOPLE & TRAINING

The people involved are willing to try AI for this use case.	YES
They can learn how to provide useful instructions and context.	YES
The proposed AI use fits reasonably into how the work is performed.	YES
Someone involved has enough subject-matter expertise to evaluate the AI output.	PARTLY
We know what quality or accuracy is acceptable for this use case.	PARTLY
People understand that AI output may need to be challenged, corrected, or rejected.	PARTLY
It is clear who remains responsible for the final work or decision.	PARTLY
People understand when human review is required.	PARTLY

RISK & SECURITY

We know whether this use is internal, employee-facing, customer-facing, vendor-facing, or public.	YES
We know who will see, use, or rely on the AI output.	YES
We understand what could happen if the AI is wrong.	YES
We know when human review is required before the output is used.	PARTLY
We have or can establish acceptable-use boundaries for this use case.	PARTLY
We can document, when appropriate, what the AI did, what information it used, who reviewed it, and what action was taken.	PARTLY

We know who is responsible for identifying and responding to security or privacy concerns.	PARTLY
--	--------

TECHNOLOGY & VENDOR

We know which technology platforms and vendors touch this process or its information.	PARTLY
We know whether those vendors use AI in relevant features or services.	PARTLY
We understand what information may be sent to or processed by the AI service.	NO
We understand how information is stored, retained, or shared.	NO
We know whether information may be used to train or improve models or services.	NO
We know what controls are available, including whether AI features can be disabled or configured.	NO
We understand who has access to the data and what happens to it if we stop using the vendor.	NO
The proposed technology is a reasonable fit for the intended business use.	NO

OPEN QUESTIONS

Which payer fields can be populated from structured EHR data without free-text PHI overcollection?	MEDIUM	OPEN
Which referral types require physician-written medical-necessity language versus nursing completion?	MEDIUM	OPEN
Can the pilot run inside an approved BAA-covered environment before portal automation?	MEDIUM	OPEN

ASSUMPTIONS

Physicians approve medical-necessity language before referral or prior-authorization submission.	UNVERIFIED
Nurses can identify missing clinical fields before a draft leaves the clinic.	UNVERIFIED
Initial pilot uses structured EHR extracts and approved templates, not payer-portal automation.	UNVERIFIED

RISKS

Draft prior-authorization rationale omits clinical detail and delays care	OPEN
LIKELIHOOD	DECISION
Medium	Reduce
PHI is sent to a tool without BAA coverage	OPEN
LIKELIHOOD	DECISION
High	Reduce

Nursing staff treat draft language as clinical judgment		OPEN
LIKELIHOOD	DECISION	
Medium	Reduce	

CONSTRAINTS

HIPAA requires BAA coverage and minimum-necessary data handling.	REGULATORY	IDENTIFIED
Physician review is required when medical necessity or diagnosis interpretation is included.	OTHER	IDENTIFIED

ACTION ITEMS

Select two high-volume referral types for structured-template pilot	PILOT PREPARATION	OPEN
Define nurse-versus-physician review rules for prior-auth drafts	GOVERNANCE	OPEN
Complete BAA and retention review for any drafting tool	INVESTIGATION	OPEN
Audit first 25 draft packets for missing fields or clinical-rationale errors	GOVERNANCE	OPEN

BUSINESS OBJECTIVES

Reduce referral packet prep time by 30% without increasing denial-related rework
EXPECTED VALUE
Time Savings
Maintain 100% physician approval for medical-necessity language in pilot packets
EXPECTED VALUE
Risk Reduction

TEAM READINESS

Name	Role	AI Familiarity	Business Knowledge	Process Knowledge	Review Capability	Risk Awareness	Notes
James Okonkwo	Clinical Operations Manager	—	—	—	—	—	—
Nina Patel	Referral Coordinator	—	—	—	—	—	—

Name	Role	AI Familiarity	Business Knowledge	Process Knowledge	Review Capability	Risk Awareness	Notes
	Lead						
Sarah Chen	Compliance & Privacy Officer	—	—	—	—	—	—

AI ROLE ASSIGNMENTS

Role	Person	Notes
AI Sponsor	Dr. Elena Vasquez	—
AI Champion	Nina Patel	—
AI Reviewer	Sarah Chen	—

OPERATING GUARDRAILS

No referral or prior-auth draft may leave the clinic without required clinical review.	HUMAN REVIEW	PROPOSED
Only BAA-covered tools may process referral packet PHI.	DATA HANDLING	PROPOSED
Pilot templates must use minimum-necessary fields by referral type.	OTHER	PROPOSED

PILOT CHARTERS

Referral and prior-authorization packets are retyped from visit data, creating nursing rework and slower specialist turnaround.	DRAFT
USERS	
Referral coordinator lead, two nursing staff, and one physician reviewer for two high-volume referral types	
IN SCOPE	
Draft referral letters and prior-authorization narratives from signed visit fields and approved payer templates	

OUT OF SCOPE

Medical-necessity decisions, payer portal automation, patient-facing communications

INFORMATION

Signed visit summary fields, diagnosis codes, medication list, payer-required administrative fields

WORKFLOW

1. Nurse selects referral type. 2. AI drafts packet from approved fields. 3. Nurse checks completeness. 4. Physician approves medical-necessity language. 5. Packet is submitted.

HUMAN REVIEW APPROACH

Nurse completeness review plus physician approval for all clinical-rationale language.

SUCCESS MEASURES

Prep time -30%; no increase in denial rework; 100% physician approval logged

RISKS SUMMARY

PHI handling, missing clinical rationale, delayed care from inaccurate drafts

CONTROLS

BAA-covered tooling; role-based access; minimum-necessary templates; compliance audit sample

OWNER

James Okonkwo

DURATION

60 days

STOP CONDITIONS

Any HIPAA incident; denial rework increases; physician review bypass occurs

DECISION POINT

Day 30: continue or narrow referral types. Day 60: expand, revise controls, or stop.

AI-Assisted Visit Summary Drafts

Sales · Completed

OPPORTUNITY BEING ASSESSED

TITLE

AI-Assisted Visit Summary Drafts

BUSINESS FUNCTION

Sales

OWNER

Dr. Elena Vasquez, Medical Director

STATUS

Ready for Pilot

VALUE PROPOSITION

Save staff time

CURRENT FRICTION

Physicians spend 15–20 minutes after each visit completing notes. Same-day completion drops during flu season backlog.

DESIRED CHANGE

Reduce documentation time to ≤10 minutes by using AI to draft structured visit summaries from EHR fields. Physicians review and sign every entry.

SUCCESS INDICATOR

Average charting time ≤10 minutes; same-day note completion ≥90%; physician satisfaction remains ≥4/5.

AI'S ROLE

AI drafts visit summaries from structured clinical fields. Physicians edit and sign before anything enters the legal chart.

ASSESSMENT

STATUS

Completed

VERSION

LIGHTNING_1.0

STARTED

9/21/2026

COMPLETED

9/21/2026

READINESS PROFILE

Dimension	Status	Key Observation
Process	READY TO INVESTIGATE	The visit-to-documentation workflow is well understood. Physicians know where AI could assist with draft visit notes. Human review before chart sign-off is already cultural.
Data & Information	SIGNIFICANT CONCERN	PHI is present in every clinical workflow. Any AI tool must have a signed BAA, defined retention policy, and audit logging. Shadow ChatGPT use for clinical content is the primary risk.
People & Training	NEEDS ATTENTION	Physicians are interested but skeptical of note accuracy. Staff need clear rules on what can enter external tools. Medical director must remain accountable for all signed chart entries.
Risk & Security	SIGNIFICANT CONCERN	Incorrect clinical summaries create patient safety and malpractice exposure. HIPAA violations from unauthorized PHI sharing must be prevented with technical and policy controls.
Technology & Vendor	NEEDS ATTENTION	EHR vendor offers limited native AI features. Ambient scribe pilot is underway but not enterprise-wide. Enterprise tool selection must precede broad rollout.

OVERALL INTERPRETATION

This opportunity is ready to investigate with compliance gates clearly defined. Clinical workflows are understood and physicians are willing to pilot. The primary constraint is PHI handling — BAA coverage and human review are non-negotiable. A 60-day pilot with three physicians at one site is the right next step once vendor security review completes.

QUESTION RESPONSES

PROCESS

We can explain how the current work is performed from beginning to end.	YES
The people involved generally perform the work in a consistent way.	YES
We understand what a good result looks like.	YES
We know where judgment, exceptions, or escalation are required.	YES
The process is documented well enough to support testing or improvement.	YES
We know who owns the process and who can approve changes to it.	YES

DATA & INFORMATION

We know what information the AI would need to perform this work.	PARTLY
We know where that information comes from and who owns it.	PARTLY
The information is accurate, complete, and usable enough for the intended purpose.	NO
We know what information would be shared with the AI and understand whether it is appropriate and permitted to share that information with the specific tool being considered.	NO
We understand whether customer, employee, financial, confidential, proprietary, regulated, or restricted information is involved.	PARTLY
We know what information should not be provided to the AI.	PARTLY

PEOPLE & TRAINING

The people involved are willing to try AI for this use case.	YES
They can learn how to provide useful instructions and context.	YES
The proposed AI use fits reasonably into how the work is performed.	YES
Someone involved has enough subject-matter expertise to evaluate the AI output.	PARTLY
We know what quality or accuracy is acceptable for this use case.	PARTLY
People understand that AI output may need to be challenged, corrected, or rejected.	PARTLY
It is clear who remains responsible for the final work or decision.	PARTLY
People understand when human review is required.	PARTLY

RISK & SECURITY

We know whether this use is internal, employee-facing, customer-facing, vendor-facing, or public.	YES
We know who will see, use, or rely on the AI output.	YES
We understand what could happen if the AI is wrong.	YES
We know when human review is required before the output is used.	PARTLY
We have or can establish acceptable-use boundaries for this use case.	PARTLY
We can document, when appropriate, what the AI did, what information it used, who reviewed it, and what action was taken.	PARTLY

We know who is responsible for identifying and responding to security or privacy concerns.	PARTLY
--	--------

TECHNOLOGY & VENDOR

We know which technology platforms and vendors touch this process or its information.	PARTLY
We know whether those vendors use AI in relevant features or services.	PARTLY
We understand what information may be sent to or processed by the AI service.	NO
We understand how information is stored, retained, or shared.	NO
We know whether information may be used to train or improve models or services.	NO
We know what controls are available, including whether AI features can be disabled or configured.	NO
We understand who has access to the data and what happens to it if we stop using the vendor.	NO
The proposed technology is a reasonable fit for the intended business use.	NO

OPEN QUESTIONS

Which ambient documentation vendors offer BAAs that meet our compliance checklist?	MEDIUM	OPEN
Can visit summaries be generated from structured EHR fields without sending full conversation transcripts to external models?	MEDIUM	OPEN
What audit trail do we need to satisfy medical director sign-off requirements?	MEDIUM	OPEN

ASSUMPTIONS

Physicians will review and edit every AI-generated note before signing the chart.	UNVERIFIED
Pilot scribe vendor BAAs cover all four clinic sites if we expand.	UNVERIFIED
Referral drafts can use structured EHR fields rather than raw conversation audio.	UNVERIFIED

RISKS

AI-generated note includes incorrect diagnosis or medication reference	OPEN
LIKELIHOOD	DECISION
Medium	Reduce
Staff enter PHI into non-BAA tools during busy clinic sessions	OPEN
LIKELIHOOD	DECISION
High	Reduce
Physicians skip review and sign AI drafts under time pressure	OPEN

LIKELIHOOD	DECISION
Medium	Reduce

CONSTRAINTS

HIPAA requires BAAs and minimum-necessary PHI handling for all AI vendors.	REGULATORY	IDENTIFIED
Payer contracts require physician attestation on all clinical documentation used for claims.	CONTRACTUAL	IDENTIFIED

ACTION ITEMS

Complete BAA and security review for ambient scribe pilot vendor	INVESTIGATION	OPEN
Define prohibited vs. permitted AI use cases in clinical policy	GOVERNANCE	OPEN
Pilot visit summary drafts with 3 physicians for 30 days	PILOT PREPARATION	OPEN
Train front-desk staff on PHI-free AI use for non-clinical templates	GOVERNANCE	OPEN

BUSINESS OBJECTIVES

Reduce average after-visit charting time from 18 minutes to ≤10 minutes within 60 days of pilot
EXPECTED VALUE
Time Savings
Maintain same-day note completion rate at ≥90% during pilot period
EXPECTED VALUE
Quality Improvement

TEAM READINESS

Name	Role	AI Familiarity	Business Knowledge	Process Knowledge	Review Capability	Risk Awareness	Notes
Dr. Elena Vasquez	Medical Director	—	—	—	—	—	—
James Okonkwo	Clinical Operations Manager	—	—	—	—	—	—

Name	Role	AI Familiarity	Business Knowledge	Process Knowledge	Review Capability	Risk Awareness	Notes
Sarah Chen	Compliance & Privacy Officer	—	—	—	—	—	—

AI ROLE ASSIGNMENTS

Role	Person	Notes
AI Sponsor	Dr. Elena Vasquez	—
AI Champion	James Okonkwo	—
AI Reviewer	Sarah Chen	—

OPERATING GUARDRAILS

No PHI may be entered into any AI tool without a signed BAA and compliance approval.	DATA HANDLING	PROPOSED
Physicians must review and edit all AI-generated clinical text before chart sign-off.	HUMAN REVIEW	PROPOSED
Pilot expansion beyond two exam rooms requires medical director written approval.	VENDOR USE	PROPOSED

PILOT CHARTERS

Physicians spend 15–20 minutes after each visit on EHR documentation, contributing to after-hours charting during peak season.	DRAFT
USERS	
3 primary care physicians at the Main Street clinic for 60-day pilot	
IN SCOPE	
Draft visit summaries from structured visit fields	
OUT OF SCOPE	

Referral letters, billing codes, prescription orders, patient-facing portal messages without physician review

INFORMATION

Structured problem list, assessment, and plan fields from EHR — no raw audio sent to external tools in phase 1

WORKFLOW

1. Physician completes visit. 2. AI drafts visit summary. 3. Physician edits in EHR. 4. Physician signs chart.

HUMAN REVIEW APPROACH

100% physician review and edit before any note is signed or referral is sent.

SUCCESS MEASURES

Charting time ≤ 10 min; same-day completion $\geq 90\%$; zero uncorrected clinical errors in audit sample

RISKS SUMMARY

Clinical accuracy; HIPAA compliance; review bypass under volume pressure

CONTROLS

BAA-required vendor; audit logging; weekly chart audit sample; compliance office escalation path

OWNER

Dr. Elena Vasquez

DURATION

60 days

STOP CONDITIONS

Any uncorrected clinical error in audit sample; HIPAA incident; physician satisfaction below 3/5 for two consecutive weeks

DECISION POINT

Day 30: continue, adjust scope, or stop. Day 60: expand, make permanent, or discontinue.