STRICTLY CONFIDENTIAL

AI Strategy and Roadmap

> *A costed AI thesis with use cases, governance, and an 18-month execution plan*

Prepared for:

[Client Confidential]

Prepared by:

Clixlogix Consulting Practice

Document Date	Q3 2025
Version	v1.0
Classification	Strictly Confidential
Distribution	Board of Directors, CEO, Executive Leadership Team

Table of Contents

Executive Summary

AI Thesis

The Client's strategic position in virtual cardiac and metabolic health rests on three durable assets that most competitors cannot replicate: 12 years of clinical outcomes data across 110,000+ completed patient programs, a multidisciplinary clinical team embedded in 32 US states, and 23 established hospital and health system partner relationships that create a recurring referral channel. The AI thesis for the Client is to convert these three assets into compounding operational and competitive advantages before the window narrows. Specifically, the Client should deploy AI to extend the clinical capacity of its NP and RD teams without proportional headcount growth, reduce between-appointment patient drop-off in the Metabolic and Hypertension programs, automate the partner reporting workflow that today consumes clinical operations bandwidth, and build a clinical outcomes intelligence capability on the proprietary dataset that no competitor can replicate because they do not have the data. This is not a thesis about becoming an AI company. It is a thesis about using AI to defend and extend a clinical delivery business that already works.

Recommended Use Cases

Use Case	Description	Phase
1. AI Care Companion	Protocol-grounded between-appointment patient support across all three program lines, with clinical escalation when warranted. Built on a new Flutter patient app, beginning with Twin Health for the Metabolic Program in Phase 1.	Phase 1-2
2. AI Session Intelligence	Rolling-transcript review during NP and RD telehealth encounters, surfacing protocol-based suggestions in real time, prompting follow-up scheduling, and producing compliance summaries for clinical management. Recommended vendor: Abridge.	Phase 1
3. Predictive Patient Engagement	Predictive engagement and churn model trained on the Client's outcomes dataset, targeted at the Metabolic and Hypertension programs, driving automated proactive outreach through existing telehealth and SMS channels.	Phase 1-2
4. AI-Assisted Partner Reporting	Automated clinical reporting layer replacing manual partner-by-partner reporting, surfacing clinical outcomes, referral conversion, and program-level ROI per partner across all 23 hospital and health system relationships.	Phase 1-2
5. Clinical Outcomes Intelligence Engine	A custom analytics capability built on the 12-year, 110,000-patient outcomes dataset, producing program-design recommendations, patient subpopulation insights, and evidence-based protocol refinements. The Client's distinctive AI asset.	Phase 2-3

Investment Envelope (18 Months)

Category	18-Month Commitment	Notes
Capital investment (platform, application, infrastructure)	$3.8M - $4.6M	Includes Flutter app build, Databricks foundation, AI application services. Figures reflect US market-rate benchmarks; Clixlogix's India-based delivery model achieves equivalent outcomes at substantially lower all-in cost.
Operating cost (model inference, vendor subscriptions, talent)	$2.1M - $2.8M annualized by month 12	Includes Abridge licensing, Twin Health pilot, and 4-6 new hires benchmarked at US market compensation rates. Clixlogix engineering partnership cost reflects offshore delivery economics.
Total program cost (18 months)	$5.9M - $7.4M	US market-rate reference range. Actual cost with Clixlogix as the AI engineering partner will be materially lower. Detailed Clixlogix-specific delivery cost plan available on request.

Expected Business Impact

The board is being asked to approve an AI investment that operates on three impact vectors simultaneously. The Hypertension Program rescue depends on predictive engagement reducing between-appointment drop-off by an estimated 18-24%, restoring the program to flat year-over-year within 12 months and returning to positive growth by month 18. The Metabolic Program re-acceleration depends on the AI Care Companion reducing passive churn, with a projected 12-15% improvement in 90-day program completion rates and a corresponding improvement in per-patient revenue. The Cardiac Program protection depends on AI session intelligence maintaining clinical quality as the program scales 40% annually without proportional NP headcount addition. Across all three programs, automated partner reporting is expected to reduce clinical operations bandwidth consumption by an estimated 30-35%, and the outcomes intelligence engine positions the Client to bring external-facing outcomes analytics to hospital partners by month 18, a capability no current competitor offers.

Build versus Buy Summary

Two use cases are recommended for buy-and-extend paths: AI Session Intelligence (Abridge as primary vendor, with Ambience as the evaluation alternate) and the Metabolic Program component of the AI Care Companion (Twin Health as the evaluation candidate). The remaining three use cases are recommended as custom builds on a Databricks data foundation, because the Client's proprietary outcomes dataset is the strategic moat and no vendor platform delivers a turnkey solution that captures

it. At US market-rate benchmarks, total vendor licensing across the 18-month program represents approximately 28-33% of total program cost; the remainder is engineering and infrastructure. With Clixlogix delivering the engineering work at offshore rates, the engineering cost component and therefore the overall program cost are both materially lower than the headline figures suggest.

Top Three Risks

- Clinical adoption resistance from NP and RD teams whose workflows will be materially augmented, particularly around AI session intelligence. Mitigation requires clinical co-design, phased rollout beginning with volunteer cohorts, and visible clinical leadership sponsorship.
- White-label patient app constraint creates a 4-6 month dependency: the Flutter app must be built before the AI Care Companion can be deployed to patients. Any delay in that platform investment delays the highest patient-impact use case.
- Vendor fit gaps not visible until evaluation, specifically the question of whether Abridge and Ambience support NP and RD telehealth encounters at production quality. If the answer is no, the session intelligence use case converts to a build path, extending the timeline by an estimated 3-4 months.

Board Decision Requested

The board is asked to approve the 18-month AI strategy and roadmap as presented, authorize an investment envelope in the range of $5.9M-$7.4M at US market-rate benchmarks (with the actual Clixlogix-delivered program cost materially lower due to offshore engineering economics), confirm the organizational commitment including the establishment of an AI Governance Committee and the sequencing of 4-6 new hires, and set the 90-day review cadence at which the executive team will report against Phase 1 milestones.

Methodology

This assessment was conducted over a six-week engagement period from July through August 2025. The work combined structured primary research with the Client's executive and clinical leadership, quantitative review of the Client's technology and data assets, and independent secondary research on the healthcare AI vendor landscape, regulatory environment, and competitive positioning of the Client's program lines.

Primary Research

The assessment team conducted 22 structured interviews across the following roles:

- Executive leadership: CEO, CFO, Chief Medical Officer, Chief Operating Officer, Head of Engineering, Head of Data Science, Head of Clinical Operations, VP of Business Development
- Clinical staff: four nurse practitioners distributed across the three program lines and three US geographic regions, three registered dietitians from the Metabolic and Cardiac programs
- Board members: two board members (roles anonymized per engagement protocol), including the lead investor representative
- Patient experience: structured review of exit interview transcripts from 60 patients who completed or discontinued programs in the prior 12 months (conducted by the Client's clinical operations team, reviewed by the assessment team)

Quantitative Review

The assessment team reviewed the following materials in structured data sessions with the Client's engineering and data science teams:

- Technology architecture: current platform stack documentation, EHR integration specifications, patient app architecture review, telehealth platform configuration
- Data infrastructure: clinical outcomes dataset schema and data quality report (Clinical outcomes dataset analysis, July 2025), data governance policies, de-identification procedures, access controls
- Operational data: program-level patient volume, completion rates, and drop-off timing by program line and patient cohort, sourced from the Client's internal business intelligence environment
- Financial data: program-level unit economics, EBITDA bridge analysis, and historical capital allocation review provided by the CFO

Secondary Research

The assessment team conducted independent research on the following:

- Healthcare AI vendor landscape: structured evaluation of ambient clinical documentation, patient engagement, predictive analytics, and healthcare data platform vendors, with particular

attention to KLAS Research 2025 Best in KLAS designations, vendor BAA posture, and deployment scale in virtual care and telehealth settings

- Competitive intelligence: review of publicly available information on the two primary competitors in the Client's program areas, including product announcements, investor communications, and market positioning
- Regulatory environment: review of HHS Office for Civil Rights HIPAA AI guidance 2025, FDA AI/ML Software as a Medical Device framework, and state-specific clinical AI regulation landscape as of Q3 2025
- Industry benchmarks: HIMSS 2025 State of Healthcare AI Survey, Rock Health 2025 Digital Health Funding Report, Menlo Ventures 2025 State of AI in Healthcare Report

Scope Exclusions

The following areas were explicitly excluded from this engagement scope and are recommended for separate workstreams:

- Revenue cycle AI: the Client's billing and coding operations involve specific payer contract nuances that require a dedicated revenue cycle assessment. Inclusion here would have lengthened the engagement by 3-4 weeks without proportional benefit to the AI strategy decision.
- Consumer marketing AI: the Client's B2B referral model means that AI-driven consumer acquisition is not relevant at this stage. This exclusion would be revisited if the Client pursues a direct-to-consumer channel in a future program line.
- AI-assisted hiring and HR operations: HR technology and workforce analytics were noted as a future opportunity but fall outside the clinical and operational AI scope defined at engagement initiation.

The AI Thesis for the Client

The Strategic Starting Point

The Client is not a technology company that happens to deliver healthcare. It is a clinical delivery organization that has built durable value through three things: the quality of its multidisciplinary clinical team, the depth of its hospital partner relationships, and 12 years of outcomes data that documents what works for which patients under which conditions. Any AI thesis that ignores these assets and pursues AI for its own sake will spend capital without building competitive advantage. Any AI thesis that fails to deploy AI at all will find those assets steadily eroded by competitors who are already moving.

The framing that guides this strategy is economic, not technological. The question is not 'what AI can we deploy?' The question is 'where does AI change the unit economics of this business in ways that are durable, defensible, and executable with the team and data the Client already has?'

That question has a clear answer. AI changes the economics of this business in four specific places: it allows more patients to be supported between appointments without proportional NP or RD time; it allows the clinical team to see more patients per clinician-hour without quality degradation; it converts the outstanding clinical reporting liability (manual, slow, partner-specific) into a competitive asset; and it converts the 12-year outcomes dataset from an underutilized archive into a live product capability. Each of these changes is directly connected to revenue, cost structure, or competitive moat.

What AI Does for This Business That Matters

Clinical Capacity Extension

The Cardiac Program is growing 40% year-over-year. The clinical team is the binding constraint on that growth. Each NP manages a patient panel whose size is limited by the time required for scheduled sessions, between-appointment responsiveness, clinical documentation, and partner reporting. AI session intelligence reduces the documentation burden per session. AI-powered between-appointment support reduces the between-appointment responsiveness burden. Taken together, these two capabilities allow the same NP to manage a materially larger patient panel without a reduction in clinical quality. The assessment team's NP interview data suggests that documentation and between-appointment administrative tasks currently consume an estimated 35-40% of productive clinical time (NP and RD interviews, n=11, August 2025). AI can reclaim a meaningful portion of that time and redirect it to patient-facing care.

Between-Appointment Engagement

The Metabolic and Hypertension programs share a structural problem: patients who disengage between appointments do not return. The drop-off is not primarily caused by clinical dissatisfaction. It is caused by the absence of meaningful between-appointment support in a program where the white-label app provides limited functionality and the clinical team's capacity for proactive outreach is constrained. A predictive engagement model that identifies patients at elevated churn risk before they disengage, combined with an AI Care Companion that provides protocol-grounded between-appointment support, addresses this problem directly. The financial impact is material: each patient retained through program completion generates approximately $2,200-$3,400 in program revenue depending on the program line

and payer mix (Client financial data, August 2025). A 15% improvement in 90-day completion rates across the Metabolic and Hypertension programs converts to a meaningful revenue recovery.

Partner Relationship Operationalization

The 23 hospital and health system partners represent both the Client's revenue engine and its most significant operational bottleneck. Each partner receives custom reports built manually by the clinical operations team. The process consumes approximately 1.5 FTE of clinical operations capacity monthly and produces reports that arrive inconsistently and vary in quality and format (Patient journey mapping with the clinical operations team, August 2025). This is a problem that frustrates existing partners and limits the Client's ability to onboard new ones, because the marginal cost of adding a 24th partner reporting relationship is high. AI-assisted partner reporting converts this from a cost center that scales linearly with partner count to a system that scales with data volume, not headcount.

Outcomes Dataset Monetization

The 12-year, 110,000-patient outcomes dataset is the Client's most underutilized asset. No competitor in the virtual cardiac and metabolic health market has a dataset of comparable depth and duration. The competitors who have launched AI features in the past six months are working from 2-4 years of data at best. The Client's dataset supports capabilities those competitors cannot yet build: subpopulation-specific protocol recommendations, long-horizon outcome predictions, and evidence-based program design refinements. The clinical outcomes intelligence engine is the use case that converts this dataset from a storage cost into a product capability. In Phase 3, it becomes the foundation for an external-facing analytics service that hospital partners will value as part of the referral relationship.

AI Moves Not Worth Making for This Business

The prioritization process surfaced approximately 18-20 candidate use cases. The following categories of AI deployment are recommended for deferral or rejection, with specific reasoning for each.

- Consumer acquisition AI: The Client's B2B referral model means that consumer-facing AI marketing does not address the actual lead generation mechanism. Hospital and health system partners generate referrals based on clinical outcomes data, relationship quality, and reporting credibility. AI can improve all three of those. AI cannot replace them with consumer advertising.
- Revenue cycle AI: Billing and coding optimization is a real opportunity but requires a separate, dedicated assessment. Including it in this strategy would dilute focus and execution capacity during Phase 1 without generating clinical or patient-facing impact. This is a Phase 3 or Year 2 workstream.
- AI-generated clinical protocols: The Client's clinical protocols are its clinical IP. Generating protocols with a general-purpose LLM without rigorous clinical review introduces patient safety risk that is not proportionate to the productivity benefit. The outcomes intelligence engine produces protocol recommendations that the clinical team reviews and adopts. That is the right architecture. Fully automated protocol generation is not appropriate for a clinical organization at this stage of AI maturity.
- Autonomous patient triage: AI that triages clinical escalations without a clinician in the loop is not appropriate for a program population that includes post-cardiac-event patients. The AI Care

Companion is designed to surface escalation recommendations to clinicians, not to make triage decisions. That design boundary is non-negotiable for this Client.

- General-purpose enterprise AI (productivity tools, meeting summarization, internal knowledge bases): These are real productivity gains but they do not change the unit economics of the clinical delivery business in the way the recommended use cases do. They can be deployed through standard enterprise software (Microsoft Copilot or equivalent) without a dedicated AI strategy workstream.

Time Horizon and Thesis Validity Conditions

This AI thesis is built for a five-year horizon from Q3 2025, with explicit recalibration triggers.

The thesis holds as long as the following conditions are true: the Client's hospital and health system partner channel remains the primary referral mechanism; commercial and Medicare Advantage payers continue to reimburse virtual specialty programs at rates that support the current margin structure; the clinical regulatory environment for AI-augmented care delivery does not shift to require FDA clearance for the specific capabilities described here; and the Client retains the clinical outcomes dataset as a proprietary asset rather than contributing it to a shared data commons.

The thesis would require material revision under any of the following conditions: a competing platform acquires the Client's outcomes dataset through a partnership or acquisition; the FDA expands its SaMD designation in a way that requires regulatory clearance for predictive engagement or session intelligence features (currently assessed as unlikely given the human-in-the-loop design, but actively monitored); or a consolidation event in the hospital partner market materially reduces the number of independent health systems generating referrals.

The specific AI move that matters most over the five-year horizon is the clinical outcomes intelligence engine. The near-term use cases (session intelligence, predictive engagement, partner reporting) create operational efficiency and patient retention improvement. The outcomes intelligence engine creates a structural competitive moat. No competitor builds that moat without the data, and the Client has 12 years of it.

The Distinctive AI Move

The Client's distinctive AI move is the conversion of 12 years of clinical outcomes data into an external-facing analytics capability that hospital partners cannot get anywhere else. This is what separates the Client's AI strategy from a generic efficiency play.

The hospital partner relationship today is a referral channel with a reporting obligation. The Client sends patients home from the hospital after a cardiac event or a diabetes diagnosis. The hospital wants to know whether those patients did well. The Client answers that question with manually produced, inconsistently delivered reports. The AI strategy converts that relationship into something more valuable: the Client becomes the organization that can tell a hospital system not only how this patient did, but what the 90th percentile outcome looks like for a patient with this profile, what the early indicators of engagement drop-off look like for this subpopulation, and what the protocol adjustments are that produced the best outcomes for patients like this one. That is a different kind of partner relationship. It is stickier, higher value, and materially harder for a new entrant to replicate.

Building this capability requires the data platform investment and the outcomes intelligence engine described in this strategy. It also requires the clinical reporting automation that ensures clean, structured data from all 23 partner relationships flowing into a unified data environment. Each of the recommended use cases feeds into the others. The strategy is a system, not a collection of independent AI experiments.

Current State and AI Readiness Assessment

The assessment team evaluated the Client's AI readiness across five dimensions: data, technology, organization, process, and capital. Each dimension was scored on a 1-5 scale following the Clixlogix Consulting Practice framework for healthcare AI readiness assessment, where 1 represents material remediation required before AI deployment and 5 represents full readiness for enterprise-scale AI programs.

Readiness Summary

Dimension	Score	Assessment Summary
Data	3.5 / 5	Strong depth; accessibility and de-identification pipeline gaps require remediation before ML workloads
Technology	2.5 / 5	White-label patient app is the primary constraint. EHR integrations and telehealth infrastructure are serviceable. Data infrastructure needs material investment.
Organization	3.0 / 5	Engineering team is capable; data science team has foundational skills but limited production ML deployment experience. Clinical leadership engagement is strong.
Process	2.5 / 5	Clinical protocols are documented but not fully digitized. Operational workflows have significant manual process dependency.
Capital	3.0 / 5	Series C runway and 8% EBITDA margin provide headroom. Board has expressed appetite. Investment sizing must be managed against EBITDA dilution during the build period.
Composite Score	2.9 / 5	The Client is AI-ready for targeted deployment on defined use cases but is not ready for broad AI platform investment without the foundational remediation steps described below.

Data Readiness - Score: 3.5 / 5

The Client's 12-year, 110,000-patient outcomes dataset is the primary AI asset and the source of the 3.5 score despite the access and pipeline gaps described below. The dataset covers three program lines at meaningful patient volume, with structured clinical outcomes data (program completion, clinical measures, episode severity, comorbidity profiles) complemented by unstructured data (session notes, patient communication transcripts, provider assessments). Approximately 65% of the dataset is in structured form amenable to direct ML use. The remaining 35% is unstructured text requiring extraction and normalization before use in training workloads (Clinical outcomes dataset analysis, July 2025).

Three specific gaps require remediation. First, the dataset currently resides in a production EHR environment with limited tooling for ML-appropriate access. There is no data lake or analytics environment that allows the data science team to run ML experiments without impacting production

systems. Second, the de-identification pipeline for HIPAA-compliant research workloads is manual and slow, creating a bottleneck for any training workflow that requires de-identified patient data. Third, data quality is uneven: the Cardiac Program data is the most complete and consistently structured; the Hypertension Program data has higher rates of missing fields and inconsistent entry across the 32-state clinical team. The quality gap in the Hypertension data is a material constraint on predictive modeling for that program line and must be addressed in parallel with the AI deployment.

Consent posture for AI use of patient data is appropriate for the operational use cases recommended here. The Client's standard consent language covers use of patient data for care improvement purposes, which the assessment team's legal review confirmed covers the intended AI applications. Research applications, specifically the clinical outcomes intelligence engine in its external-facing form, may require supplemental consent for a cohort of historical patients depending on state-specific requirements. The legal review of consent posture for the outcomes intelligence engine is flagged as a required Phase 2 gate.

Technology Readiness - Score: 2.5 / 5

The white-label patient app is the most significant technology constraint in the assessment. The current platform cannot be extended with new features. The AI Care Companion, which is the highest patient-impact use case, cannot be deployed until a new Flutter or React Native patient app is built on top of the white-label platform's API layer. This dependency is the single longest-lead-time item in the 18-month roadmap, and Phase 1 begins with the Flutter app build for that reason. The app build is estimated at 3-4 months from authorization to production deployment, which means the AI Care Companion reaches patients in month 4-5 at the earliest.

The EHR integrations are functional across the three primary EHR systems used by the Client's hospital partners, but integration depth varies. Some partner integrations support bidirectional data flow; others are read-only. The clinical reporting automation use case requires standardized data extraction from all 23 partner EHR environments, which will require integration remediation for approximately 7-8 partners with read-only or non-standard integration configurations (Technology audit, August 2025).

The telehealth infrastructure supports the AI session intelligence use case with modifications. The current platform captures session video and audio but does not route audio to a real-time transcription layer. Integration of Abridge or an equivalent ambient documentation platform requires API-level integration with the telehealth platform. The Client's Head of Engineering assessed this integration as a 6-8 week effort.

There is no current data infrastructure (data lake, feature store, vector store, or ML pipeline tooling) beyond the production EHR environment and a basic business intelligence layer. The Databricks Lakehouse investment recommended in this strategy creates the foundational data infrastructure that supports three of the five recommended use cases.

Organizational Readiness - Score: 3.0 / 5

The engineering team has 14 engineers with mixed full-stack and backend specializations. The team has built and maintained the current platform, the EHR integrations, and the telehealth infrastructure. AI engineering experience is limited: one engineer has production ML deployment experience; the others

have exposure to AI through side projects or coursework but have not deployed AI to production at scale. The team is capable and motivated; the capability gap is specific to AI engineering discipline and can be closed through a combination of targeted hiring and partnership with an AI engineering practice.

The data science team has 3 analysts and one senior data scientist. The senior data scientist led the predictive cardiac risk model that produced useful signals but was not operationalized. This experience is informative: the team can build predictive models, but the gap between a model that produces signals and a model that is integrated into a clinical workflow is an engineering and process problem that the current team has not yet solved. The recommended organizational model addresses this gap explicitly.

Clinical leadership engagement with AI is stronger than the assessment team expected at engagement initiation. The CEO's clinical background as a former cardiologist means that the AI-augmentation framing (AI that extends clinical judgment rather than replacing it) resonates at the leadership level. The Chief Medical Officer has been actively engaged in the AI strategy design and is prepared to lead clinical co-design of the AI Care Companion and session intelligence features. This is a meaningful organizational strength.

NP and RD sentiment toward AI augmentation is mixed. Interview data surfaced two distinct cohorts: a group of early adopters (approximately 40% of interviewees) who are enthusiastic about AI session intelligence specifically, and a group of skeptics (approximately 35% of interviewees) who express concern about AI affecting the quality of the clinician-patient relationship. The remaining 25% express general openness contingent on the design of the AI features. The change management plan in Section 9 addresses this distribution directly.

Process Readiness - Score: 2.5 / 5

Clinical protocols for all three program lines are documented and maintained by the clinical operations team. The Cardiac Program protocols are the most complete and have been updated most recently. The Hypertension Program protocols have not been updated in 18 months and have several outstanding review items flagged by the CMO. The AI Care Companion and session intelligence features depend on machine-readable protocol representations; converting the current document-based protocols to a format usable for AI prompting and real-time clinical decision support is a 4-6 week effort estimated by the data science team.

Operational workflows for partner reporting, patient onboarding, and between-appointment outreach are partially documented but heavily dependent on individual staff knowledge and manual execution. The partner reporting process in particular operates through a combination of spreadsheets, email threads, and direct EHR access that varies by partner. This process documentation gap is the primary obstacle to clean, automated reporting; the AI-assisted partner reporting use case cannot be deployed until the underlying reporting workflows are documented, standardized, and mapped to the data available in each partner's EHR integration.

Capital Readiness - Score: 3.0 / 5

The Client's 8% EBITDA margin on approximately $45M in revenue generates approximately $3.6M in EBITDA annually. The Series C raise completed 14 months prior to this engagement provides runway beyond the 18-month program window. The board has expressed explicit appetite for AI investment and

has asked for a defensible investment thesis rather than a request to defer. This is a favorable capital position relative to most organizations at comparable scale.

The principal capital readiness constraint is EBITDA dilution management. The investment envelope of $5.9M-$7.4M over 18 months referenced in the executive summary reflects US market-rate benchmarks for comparable engineering talent and represents a material commitment relative to annual EBITDA at those rates. With Clixlogix delivering the engineering and build components at offshore rates, the actual program cost is materially lower, which improves the EBITDA dilution profile during the build period. The phased structure of the roadmap is designed to generate measurable business impact from Phase 1 investments before the full Phase 2 capital commitment is made, providing the board with evidence-based recalibration points at months 6 and 12. The alternative of front-loading capital without phased validation would expose the investment to full execution risk without staged off-ramps.

Use Case Prioritization

Candidate Use Case Generation

Candidate use cases were generated through four parallel inputs: clinical workshops with the NP and RD teams across the three programs; structured interviews with board members and the executive leadership team; review of the two prior AI experiments (the chatbot pilot and the predictive cardiac risk model); and competitive analysis of AI features deployed by competitors in the prior six months. The initial candidate set was further informed by the assessment team's secondary research on healthcare AI vendor capabilities as of Q3 2025.

The generation process produced approximately 20 distinct candidate use cases organized across four categories: patient-facing AI (what patients experience directly), clinician-facing AI (what NPs and RDs experience in their workflow), operational AI (what clinical and administrative operations teams experience), and commercial AI (what business development and partner management teams experience). The full candidate list is reproduced in the prioritization matrix below.

Prioritization Framework

The assessment team scored each candidate use case on a weighted four-dimension framework adapted from the Clixlogix Consulting Practice healthcare AI prioritization model:

Dimension	Weight	What It Measures
Expected Business Impact	35%	Revenue uplift, cost reduction, retention improvement, or competitive moat creation, scored against the specific business context of this Client
Feasibility Given Current Readiness	30%	Data availability and quality, technology integration complexity, organizational capability match, and regulatory risk, scored against the readiness assessment in Section 4
Strategic Fit with AI Thesis	25%	Degree to which the use case advances the four thesis pillars: clinical capacity extension, between-appointment engagement, partner relationship operationalization, or outcomes dataset monetization
Time-to-Value	10%	Speed to first measurable business impact, reflecting the board's 90-day reporting commitment and the competitive pressure in the Metabolic and Hypertension program areas

Full Candidate Use Case Ranking

Use Case	Category	Impact	Feasibility	Fit	Time	Score	Status
AI Care Companion (Metabolic + Cardiac)	Patient-Facing	5	3.5	5	3	4.1	**Recommended**
AI Session Intelligence (NP + RD)	Clinician-Facing	4.5	4	5	4	4.3	**Recommended**
Predictive Patient Engagement + Churn	Operational	4.5	4	4.5	4	4.3	**Recommended**
AI-Assisted Partner Reporting	Operational	4	3.5	4.5	3.5	3.9	**Recommended**
Clinical Outcomes Intelligence Engine	Operational	5	3	5	2	3.8	**Recommended**
Hypertension Program Protocol AI	Clinician-Facing	3	2.5	3.5	2.5	2.9	Deferred
AI-Powered Intake and Onboarding	Patient-Facing	3	3	3	3.5	3.0	Deferred - Phase 3
Automated Prior Authorization Support	Operational	3.5	2	3	2	2.7	Deferred
NP Panel Optimization Engine	Operational	3	2.5	3	2.5	2.8	Deferred - Phase 3
AI-Assisted Clinical Documentation (Non-Session)	Clinician-Facing	3	3.5	3	3.5	3.2	Subsumes into Use Case 2
Patient Sentiment and NPS Prediction	Patient-Facing	2.5	3	2.5	3	2.7	Deferred
Referral Conversion AI (Partner-Facing)	Commercial	3.5	2	3	2	2.6	Deferred
AI-Generated Patient Education Content	Patient-Facing	2	4	2	4	2.6	Rejected - Low Fit
AI Chatbot (General Health Questions)	Patient-Facing	1.5	3.5	1.5	3.5	2.2	Rejected - Prior Failure
Revenue Cycle AI (Coding + Billing)	Operational	3.5	2	2.5	1.5	2.5	Out of Scope
HR and Workforce Analytics AI	Operational	2	3	1.5	3	2.2	Out of Scope

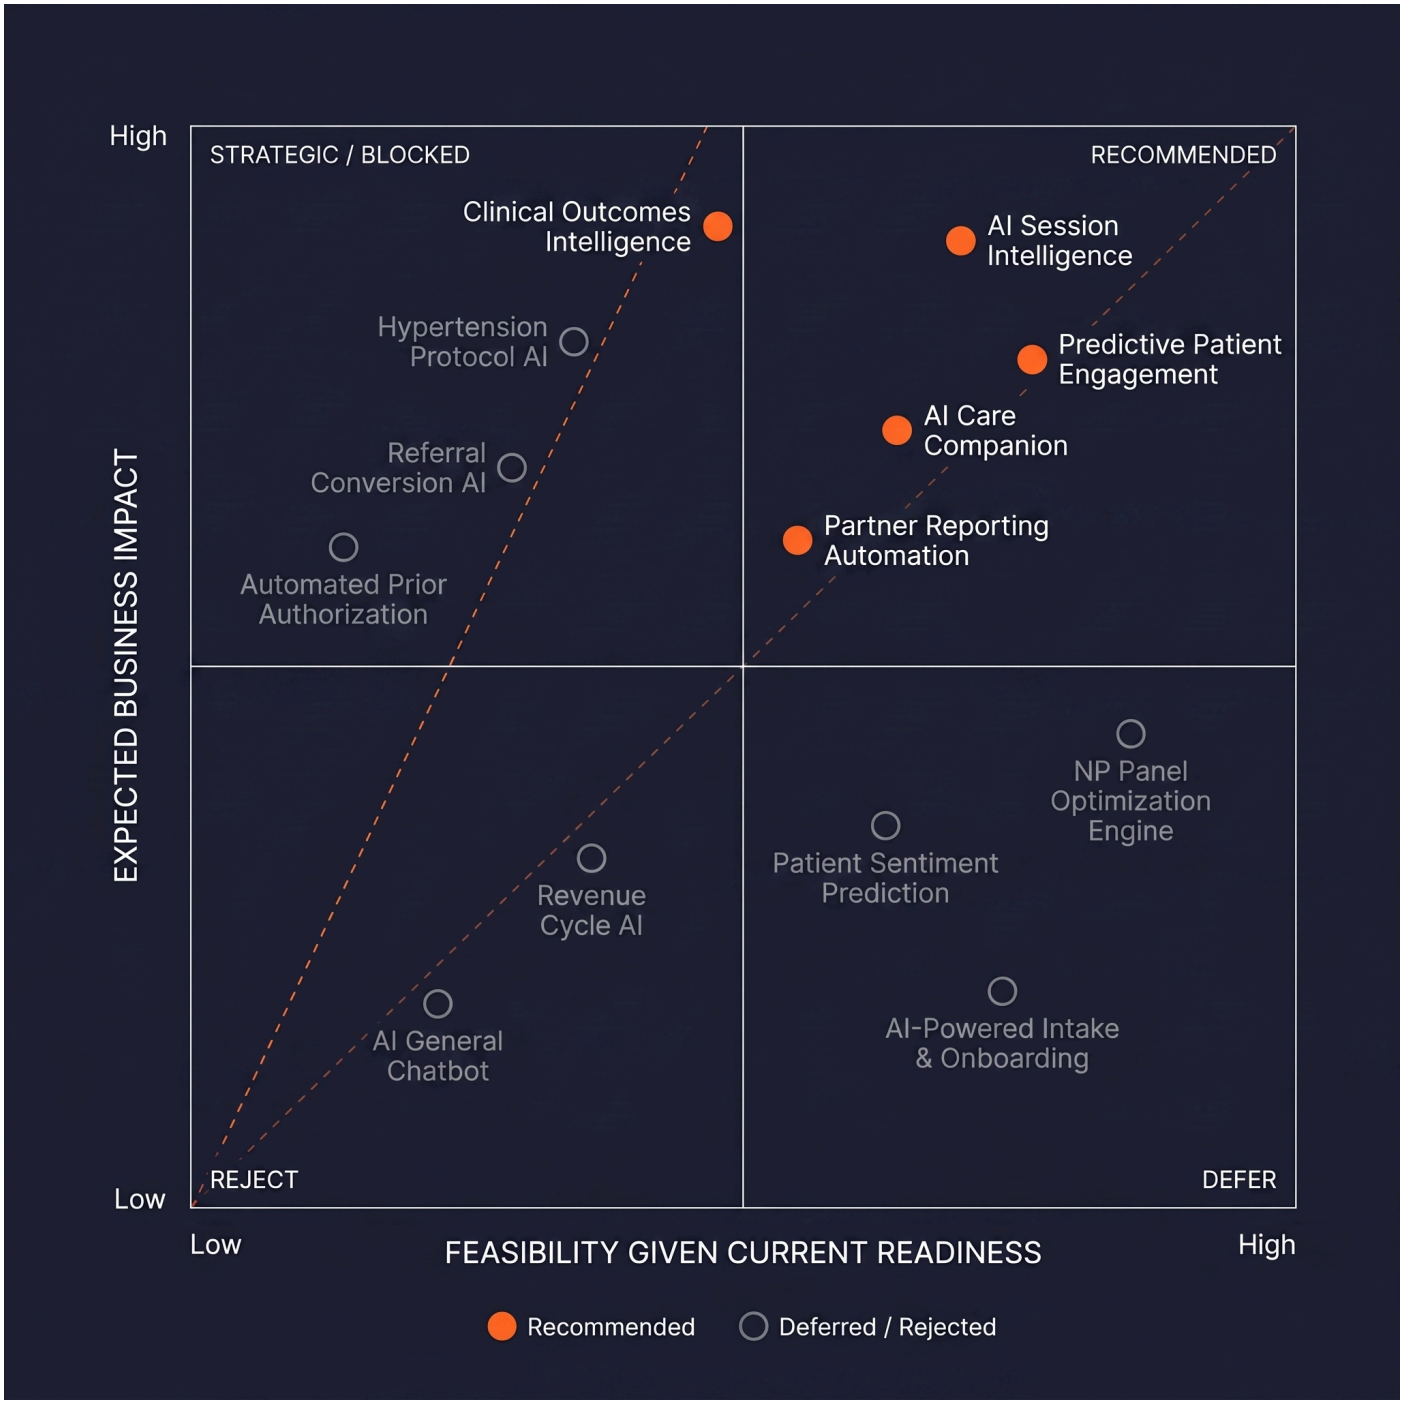
High
STRATEGIC / BLOCKED
RECOMMENDED
Clinical Outcomes Intelligence
AI Session Intelligence
Hypertension Protocol AI
Predictive Patient Engagement
Referral Conversion AI
AI Care Companion
Automated Prior Authorization
Partner Reporting Automation
EXPECTED BUSINESS IMPACT
NP Panel Optimization Engine
Patient Sentiment Prediction
Revenue Cycle AI
AI General Chatbot
AI-Powered Intake & Onboarding
REJECT
DEFER
Low
Low
FEASIBILITY GIVEN CURRENT READINESS
High
Recommended
Deferred / Rejected

Recommended Use Cases - Detailed Treatment

Use Case 1: AI Care Companion

The clinical problem this use case addresses is specific and documented. Patients in the Metabolic and Hypertension programs who disengage between appointments rarely re-engage. The current white-label app provides medication reminders and basic health logging but no meaningful clinical interaction between scheduled sessions. For a patient managing Type 2 diabetes or chronic hypertension, the three to seven days between appointments are when behavior slips, medications are missed, and the conditions worsen quietly. The clinical team does not have capacity to proactively engage every at-risk patient between sessions. The AI Care Companion fills that gap.

The AI approach is a clinically grounded conversational AI layer, anchored to the Client's specific clinical protocols for each program line, that provides protocol-appropriate responses to between-appointment patient questions, surfaces evidence-based behavioral nudges based on patient-reported data, and identifies clinical escalation signals that warrant notification to the assigned NP or RD. The critical design constraint is that the system does not diagnose, does not independently modify care plans, and does not make autonomous clinical decisions. It surfaces information and escalations to clinicians; clinicians act on them.

The business impact estimate is grounded in two assumptions: a 12-15% improvement in 90-day program completion rates in the Metabolic Program, and a 15-20% reduction in unscheduled drop-off in the Hypertension Program. At the Client's current patient volume and per-patient economics, these improvements generate an estimated $2.8M-$4.1M in incremental annual revenue by month 18, with full program ramp. The basis for these estimates is a combination of published outcomes from similar patient engagement interventions (Menlo Ventures 2025 State of AI in Healthcare Report) and the Client's own historical data on the revenue impact of program completion rate improvements (Client financial data, August 2025).

Infrastructure required: the Flutter patient app (the platform prerequisite), a BAA-governed commercial LLM API for the conversational layer, a vector store for clinical protocol retrieval, and integration with the existing EHR for patient context. Build-versus-buy stance: buy-and-extend hybrid, beginning with Twin Health for the Metabolic Program and building the custom conversational layer in parallel for full deployment in Phase 2. The signal that this use case is working is 90-day program completion rate by program line, tracked monthly from first patient deployment.

Use Case 2: AI Session Intelligence

The clinical problem is documentation burden. An NP conducting a 30-minute telehealth session with a Cardiac Program patient spends an estimated 12-18 minutes on post-session documentation: clinical notes, care plan updates, follow-up scheduling, compliance documentation, and partner reporting data entry (NP and RD interviews, n=11, August 2025). That documentation burden is the reason that the same NP cannot see more patients per day without extending work hours or sacrificing documentation quality. AI session intelligence addresses the burden directly.

The AI approach is ambient documentation: a real-time transcription and analysis layer running during the NP or RD session that produces a structured clinical note draft, identifies care plan items mentioned during the session, surfaces protocol-appropriate suggestions based on what the patient reports, and generates the compliance documentation fields required for partner reporting. The clinician reviews and approves the AI output; the system does not write to the EHR autonomously. The human-in-the-loop design is not optional for this Client. It is required by the clinical safety governance framework and by the standard of care for NP-supervised chronic disease management programs.

The expected business impact is capacity expansion without proportional headcount cost. If AI session intelligence reduces post-session documentation time from an average of 15 minutes to 5 minutes per session, and each NP conducts an average of 8 sessions per day, the capacity unlocked is equivalent to approximately 1.3 additional sessions per NP per day. Across 80 NPs, that unlocked capacity is equivalent to approximately 100 patient sessions per day without a new hire. At the Cardiac Program's 40% annual growth rate, that capacity buffer is not a convenience; it is the difference between managing growth and being forced into reactive hiring.

Build-versus-buy stance: buy. Abridge is the recommended primary evaluation candidate. The risks and evaluation criteria are addressed in detail in Section 6. The metric that signals this use case is working is post-session documentation time per NP, measured by comparing session end time to chart close time in the EHR before and after deployment.

Use Case 3: Predictive Patient Engagement and Churn Modeling

The clinical problem is visibility. The current clinical workflow is reactive: a patient misses an appointment, the care coordinator follows up, the patient has already decided to disengage. The predictive engagement model changes the workflow to proactive: the model identifies patients whose behavioral and clinical signals suggest elevated disengagement risk in the next 14 days, and the system routes those patients to targeted outreach through the existing telehealth and SMS channels before the appointment is missed. This is not a new idea in population health management, but the Client's 12-year outcomes dataset is an unusually strong foundation for building a model that is specific to this patient population and this program structure.

The AI approach is a supervised predictive model trained on historical patient engagement and outcome data, producing a daily engagement risk score for each active patient. Features include appointment adherence history, in-app engagement patterns (currently limited by the white-label app but improving with the Flutter app), patient-reported outcome scores, clinical measures trends, and program-line-specific behavioral patterns identified in the historical dataset. The model output feeds an automated outreach workflow that routes the highest-risk patients to proactive NP or care coordinator contact.

The expected business impact is retention improvement in the two struggling programs. The Metabolic Program currently has a 60-day drop-off rate of approximately 22%; the Hypertension Program has a 60-day drop-off rate of approximately 34% (Client financial data, August 2025). A predictive model that catches 40-50% of imminent churn events and converts half of those to retained patients through proactive outreach would reduce the Metabolic drop-off rate by approximately 4-5 percentage points and the Hypertension drop-off rate by approximately 6-8 percentage points. At current patient volumes and per-patient economics, that retention improvement generates an estimated $1.6M-$2.4M in annual revenue recovery.

Build-versus-buy stance: build with data platform foundation. The Client's dataset is the strategic advantage here. Vendor platforms provide general-purpose risk models; a custom model trained on this specific patient population will outperform them. The signal that this use case is working is 30-day and 60-day retention rates by program line, compared to historical baseline.

Use Case 4: AI-Assisted Clinical Partner Reporting Automation

The operational problem is structural. Each of the 23 hospital and health system partners requires a custom monthly report that documents clinical outcomes for their referred patients, referral conversion rates, and program ROI for the partnership relationship. The current process is entirely manual: clinical

operations staff extract data from the EHR, format it according to each partner's specific requirements, and distribute via email. The process takes approximately 1.5 FTE of clinical operations capacity per month, produces reports that arrive an average of 8-12 business days after month close, and creates inconsistency in how outcomes are presented across partners.

The AI approach is a three-layer automation: a standardized data extraction and normalization layer that pulls structured outcomes data from all 23 partner EHR integrations into a unified reporting environment; a natural language generation layer that produces narrative report sections using a BAA-governed LLM; and a partner-specific formatting layer that applies each partner's preferred report structure. The clinical operations team reviews and approves the generated reports before distribution; the system does not send partner communications autonomously.

The expected business impact is two-dimensional. First, the 1.5 FTE of clinical operations capacity currently consumed by partner reporting is redirected to patient-facing clinical coordination, equivalent to an estimated $145,000-$175,000 in annual productivity recovery. Second, the quality and consistency improvement in partner reporting is expected to strengthen partner retention and accelerate new partner onboarding. The assessment team estimated conservatively that 2-3 new hospital partner relationships could be onboarded in the 12 months following deployment of automated reporting, on the basis that the current reporting bottleneck is one of the named reasons prospective partners cite for delayed program adoption (Business development interviews, August 2025). At an average annual partner revenue of $1.2M-$1.8M per health system relationship, the partner acquisition impact dwarfs the operational savings.

Build-versus-buy stance: build with data platform foundation. The Client's 23 unique reporting relationships each have custom configurations that make a vendor platform approach as complex as a build. The signal that this use case is working is reporting cycle time (from month close to partner delivery) and partner satisfaction scores, collected quarterly.

Use Case 5: Clinical Outcomes Intelligence Engine

The strategic problem is that the Client's most valuable asset is not being used. The 12-year, 110,000-patient outcomes dataset documents what works for which patients, under which clinical protocols, in which program structures. The data science team has not had the infrastructure or the organizational bandwidth to build the analytical capability that converts this dataset into program design intelligence. The clinical outcomes intelligence engine is that capability.

The AI approach is a custom analytics and machine learning platform built on Databricks Lakehouse infrastructure. In its Phase 2 form, the engine produces internal program-design insights: which protocol variants produce the best outcomes for which patient subpopulations, where the Hypertension Program's clinical engagement falls short relative to the Cardiac Program, which early clinical indicators are most predictive of long-term program success across program lines. These insights inform protocol refinements and clinical team training that improve program outcomes at the aggregate level. In its Phase 3 form, the engine produces external partner-facing analytics: hospital partners receive not just reports on their referred patients, but benchmarked insights on how their patient population compares to the full dataset and what the evidence-based protocol adjustments are that could improve outcomes.

The expected business impact is compounding and long-term. In Phase 2, internal program design improvements are estimated to produce a 2-4% improvement in program completion rates across all three programs by month 18, generating an estimated $1.1M-$2.2M in incremental annual revenue at scale. In Phase 3, external partner analytics is expected to function as a material driver of partner retention and new partner acquisition, with an estimated 15-20% increase in partner renewal

commitment intensity (a qualitative estimate based on business development interview data). The revenue impact of partner retention is significant but difficult to model with precision at this stage; the recommendation is to treat the Phase 3 partner analytics capability as a strategic option with high expected value rather than a modeled revenue line.

Build-versus-buy stance: build on Databricks foundation. No vendor platform delivers a turnkey clinical outcomes intelligence engine for this specific patient population and program structure. The signal that this use case is working is program completion rate improvement by program line and protocol adherence rates by clinical subpopulation, measured quarterly.

Deferred Use Cases and Rationale

Use Case	Deferral Reason	Revisit Trigger
Hypertension Program Protocol AI	Hypertension protocols have not been updated in 18 months and have outstanding clinical review items. Building AI on top of outdated protocols would encode the existing gaps. Revisit after protocol refresh, targeted for Phase 2.	Protocol refresh completed and CMO approval secured
AI-Powered Intake and Onboarding	High feasibility but lower strategic priority than the retention-focused use cases. Patient acquisition is not the binding constraint at current growth rates. Phase 3 candidate.	Hypertension rescue metrics met; Cardiac growth plateaus
Automated Prior Authorization Support	Data readiness for payer-specific authorization workflows is insufficient. Payer integration complexity requires a dedicated workstream. Year 2 recommendation.	Databricks foundation live; payer API integration scoped
NP Panel Optimization Engine	Organizational capacity for change management during Phase 1 is finite. NP panel optimization introduces workflow change for the entire clinical team simultaneously with session intelligence deployment. Phase 3 candidate.	Session intelligence fully deployed and adoption stable
Referral Conversion AI	The referral conversion problem is a relationship and reporting problem, not a data problem. Partner reporting automation (Use Case 4) addresses the root cause. Revisit after 6 months of automated reporting data.	12+ months of structured partner reporting data available
AI General Chatbot	The Client's prior chatbot pilot was sunset after 6 months due to low engagement. The failure analysis attributed the underperformance to lack of clinical grounding. The AI Care Companion addresses this with protocol-grounded design. A general chatbot is not recommended.	Explicitly rejected, not revisited

Build versus Buy versus Partner

> *This section names specific vendor platforms as recommended evaluation candidates for each use case. The healthcare AI market is moving rapidly. Platforms named here reflect the market landscape as of Q3 2025. Specific selections require current-state evaluation against the Client's criteria during the implementation phase, including BAA terms, integration with the Client's existing platform stack, support for NP and RD telehealth encounters where applicable, pricing at the Client's scale, vendor financial stability, and product roadmap alignment. The named platforms are starting points for evaluation, not final selections.*

Use Case 1: AI Care Companion

Dimension	Assessment
Recommended path	Buy-and-extend hybrid: Twin Health for the Metabolic Program in Phase 1; custom Flutter app with proprietary clinical protocol AI layer for full deployment in Phase 2
Build description	Custom Flutter or React Native patient app on top of the white-label platform's API, with a BAA-governed commercial LLM (OpenAI Healthcare BAA, Anthropic Enterprise, or Google Cloud Healthcare API) providing the conversational AI layer, anchored to the Client's clinical protocols via a vector retrieval architecture
Build time-to-value	Phase 1 app: 3-4 months. Phase 2 AI Care Companion full deployment: months 6-9 from authorization

Twin Health - Primary Metabolic Program Evaluation Candidate

Market positioning: AI Digital Twin pioneer for metabolic health; published randomized controlled trial results in the New England Journal of Medicine Catalyst; outcomes validated by the American Heart Association, American Diabetes Association, and American College of Cardiology.

Twin Health's Whole Body Digital Twin platform is the first metabolic health vendor to publish a gold-standard randomized controlled trial in the New England Journal of Medicine Catalyst, with peer-reviewed results also published in journals of the American Heart Association, American Diabetes Association, and American College of Cardiology (Twin Health published clinical outcomes, 2025). In August 2025, Twin Health closed a $53M investment and announced validation by the Validation Institute confirming that its platform outperforms legacy digital health solutions for Type 2 diabetes, pre-diabetes, and obesity on clinical and financial outcomes. The Validation Institute's certification is specifically relevant for commercial insurer and Medicare Advantage payer conversations, which constitute the Client's primary payer mix. Twin Health operates under a BAA and has production deployments in US commercial and Medicare Advantage employer and health plan environments. The evaluation question for Twin Health remains the same protocol extension question that governs all AI Care Companion vendor evaluations: whether the platform's patient engagement architecture can be extended with the Client's proprietary clinical protocols, or whether it operates as a protocol replacement. Twin Health's outcomes-based pricing model -- where clients pay when members achieve specific clinical outcomes such as A1C reduction or medication elimination -- is also worth evaluating

against the Client's financial model, as it creates a shared-risk structure that reduces upfront vendor cost at the expense of per-outcome payment obligations.

Story Health by Innovaccer - Secondary Cardiac Program Evaluation Candidate

Market positioning: AI-powered continuous cardiovascular specialty care platform (heart failure, hypertension, AFib); acquired by Innovaccer in September 2025 and now embedded in Innovaccer's Healthcare Intelligence Cloud alongside Gravity, Atlas, and Cured.

Story Health was acquired by Innovaccer in September 2025, becoming the specialty care delivery arm of Innovaccer's Healthcare Intelligence Cloud. Prior to acquisition, Story Health had established validated cardiovascular care deployments at ChristianaCare, Intermountain Health, and Saint Luke's Mid America Heart Institute, with published outcomes including a 6.9% 30-day heart failure all-cause readmission rate against the national average of 18.1%. The acquisition is directly relevant to the Client's vendor evaluation because selecting Story Health now means evaluating the full Innovaccer platform stack. If the Client pursues the platform-led path for Use Cases 3 and 4 (predictive engagement and partner reporting via Innovaccer Gravity and Atlas), Story Health by Innovaccer for the Cardiac Program component of Use Case 1 becomes a natural extension of that single vendor relationship rather than a separate contract and integration. This platform consolidation scenario -- Innovaccer covering cardiac specialty care, predictive analytics, and partner reporting -- should be modeled as an explicit alternative to the build-first recommendation in Phase 1 evaluation. The diabetes and COPD program tracks planned by Story Health post-acquisition had not launched as of Q3 2025, which limits its current applicability to the Metabolic Program. The two-platform scenario (Twin Health for Metabolic, Story Health by Innovaccer for Cardiac) and the consolidated Innovaccer platform scenario should both be evaluated during Phase 1 against the build alternative.

Use Case 2: AI Session Intelligence

Dimension	Assessment
Recommended path	Buy. Abridge as primary evaluation candidate, Ambience as alternate. Build is the fallback if NP and RD telehealth encounter support fails vendor evaluation.
Critical evaluation criterion	Both platforms are predominantly built for physician encounters in ambulatory and inpatient settings. The Client uses NPs and RDs in telehealth sessions. NP and RD telehealth encounter support must be validated at production quality before vendor selection. This is a gating question.
Build alternative	Custom LLM-based session intelligence on a healthcare-BAA-governed commercial LLM with a specialty-specific prompting layer trained on the Client's clinical protocols. Adds approximately 3-4 months to time-to-value.

Abridge - Primary Evaluation Candidate

Market positioning: 2025 and 2026 Best in KLAS winner for Ambient AI; deployed at 200+ health systems including UPMC, Mayo Clinic, Johns Hopkins, and Duke Health.

Abridge holds consecutive Best in KLAS designations for the Ambient AI category in 2025 and 2026, and has the largest enterprise-scale deployment footprint of any ambient documentation platform as of

Q3 2025, with deployments confirmed at more than 200 health systems (KLAS Research 2025 Best in KLAS Awards; Abridge 2025 deployment data). The platform's specialty-specific clinical terminology recognition is materially more accurate than general-purpose transcription for the chronic disease management specialties relevant to the Client's program lines. Abridge's BAA posture and HIPAA compliance architecture have been validated at health system scale. The evaluation criterion that distinguishes Abridge's fit for this Client is NP and RD telehealth encounter support: the platform was built primarily for physician encounters, and the Client's clinical team is NP and RD-led in telehealth settings. The evaluation team must confirm that Abridge's entity recognition and note-generation logic performs at production quality for these encounter types before proceeding to deployment.

Ambience - Alternate Evaluation Candidate

Market positioning: $1.25B valuation (Series C, July 2025) ambient AI platform for documentation, coding, and clinical documentation integrity; selected by Cleveland Clinic after a six-month competitive evaluation.

Ambience reached unicorn status in July 2025 with a $243M Series C raise at a $1.25 billion valuation, positioning it as the second-highest valued company in the ambient AI space behind Abridge (valued at approximately $5.3 billion at Series E). Ambience's platform has broadened beyond ambient documentation to cover clinical documentation integrity and point-of-care coding, which gives it a wider footprint in revenue cycle workflows than a pure ambient scribe. Cleveland Clinic selected Ambience after a six-month competitive pilot that included five ambient AI vendors, which is the most directly comparable public reference evaluation to the Client's use case. Ambience has been growing in complex subspecialty and inpatient settings, which may or may not match the Client's telehealth-first encounter model. The mid-market pricing may be more favorable than Abridge's enterprise pricing at the Client's scale of approximately 120 NP and RD users. The same NP and RD telehealth encounter validation criterion applies. If both platforms fail the telehealth NP/RD encounter validation, the build alternative is recommended.

Use Case 3: Predictive Patient Engagement and Churn Modeling

Dimension	Assessment
Recommended path	Build with data platform foundation. Databricks or Snowflake Healthcare Data Cloud as the data foundation; custom predictive model built on the Client's proprietary outcomes dataset.
Build rationale	Vendor platforms provide general-purpose risk models trained on broad population datasets. The Client's 12-year outcomes dataset is the competitive advantage; a custom model trained on this specific patient population will outperform a general model on the metrics that matter for this use case.
Estimated build timeline	4-6 months from data infrastructure availability to Phase 1 deployment in the Metabolic Program.

Innovaccer - Platform-Led Alternative

Market positioning: Healthcare Intelligence Cloud covering data unification (Gravity), population health analytics (Atlas), patient engagement (Cured), and continuous specialty care delivery (Story Health by Innovaccer); $275M

Series F raised January 2025; Best in KLAS across Data Analytics for Providers, Payers, and CRM categories in 2026.

With the Story Health acquisition completed in September 2025, Innovaccer now offers the broadest consolidated platform footprint of any vendor relevant to this Client's use cases: Gravity covers data unification and AI orchestration, Atlas covers population health analytics and partner reporting, Cured covers patient engagement and communication, and Story Health by Innovaccer covers continuous cardiovascular specialty care delivery. This means a decision to select Innovaccer for Use Cases 3 and 4 now also creates a natural path to evaluating Story Health by Innovaccer for the Cardiac Program component of Use Case 1, without adding a new vendor relationship. The platform consolidation argument across Use Cases 1, 3, and 4 simultaneously is materially stronger post-acquisition than it was before September 2025. The trade-off assessment remains the same: the Client's 12-year proprietary outcomes dataset is the strategic moat, and Innovaccer's platform would ingest that dataset into a vendor-controlled environment where predictive models are trained on Innovaccer's methodology rather than built natively on the Client's data. The assessment team's view is that the strategic value of owning the model outweighs the operational convenience of platform consolidation. However, if Phase 1 engineering capacity is more constrained than projected, the full Innovaccer platform path covering Use Cases 1, 3, and 4 is the most credible single-vendor alternative.

Health Catalyst - Alternative for EHR-Integration-Heavy Environments

Market positioning: Population health analytics platform with mature predictive risk modeling and strong EHR integration depth.

Health Catalyst is the stronger evaluation candidate if the Client's EHR integration complexity exceeds the assessment team's current estimate. Health Catalyst has mature integration tooling for the major EHR systems used by the Client's hospital partners and a predictive risk modeling layer that can be configured to the Client's program population. The evaluation criterion is whether Health Catalyst's predictive model configuration is flexible enough to leverage the Client's proprietary historical dataset as the primary training source, or whether it defaults to Health Catalyst's aggregate benchmarking dataset. If the former is achievable, Health Catalyst is a viable alternative to the custom build.

Use Case 4: AI-Assisted Clinical Partner Reporting Automation

Dimension	Assessment
Recommended path	Build with data platform foundation. Custom reporting service on Databricks or Snowflake Healthcare Data Cloud with BAA-governed LLM for narrative generation.
Build rationale	The Client's 23 partner reporting relationships each have custom configurations. Any vendor platform requires substantial customization regardless; the build alternative produces an owned asset for comparable implementation effort.
Estimated build timeline	3-4 months for MVP serving the 4 largest partners; extended to all 23 partners through Phase 2.

Innovaccer Atlas - Platform Alternative

Market positioning: Population health reporting and partner analytics platform; part of the Innovaccer Healthcare Intelligence Cloud; $275M Series F (January 2025); 2026 Best in KLAS for Data Analytics for Providers.

Innovaccer Atlas is recommended as the vendor platform alternative if Innovaccer is also selected for Use Case 3, because platform consolidation across both use cases -- and potentially Use Case 1 via Story Health by Innovaccer for the Cardiac Program -- eliminates multiple data integration layers and reduces vendor management complexity to a single contract. Atlas's reporting and analytics capabilities are built for partner-facing population health reporting at health system scale. Innovaccer's financial position is strong: the $275M Series F closed in January 2025 with participation from Kaiser Permanente and Banner Health as strategic investors, which reduces vendor financial instability risk materially. The evaluation question for Atlas is customization depth: the Client's 23 partner reporting relationships each have distinct requirements, and the evaluation must confirm that Atlas supports per-partner report configuration without professional services dependency for each change.

Arcadia - Reporting-First Alternative

Market positioning: Population health analytics platform with a reporting-first architecture; alternative to Innovaccer if broader platform consolidation is not selected.

Arcadia is the appropriate evaluation candidate if the Client does not select Innovaccer for Use Case 3 and prefers to evaluate a reporting-focused platform for Use Case 4 independently. Arcadia's architecture is designed for population health reporting to health system stakeholders, which maps directly to the partner reporting use case. The evaluation criteria are identical to Innovaccer Atlas: per-partner customization capability, BAA posture, and integration with the Client's existing EHR connections.

Use Case 5: Clinical Outcomes Intelligence Engine

Dimension	Assessment
Recommended path	Build on Databricks Lakehouse foundation. Custom model architecture trained on the Client's proprietary dataset. External Tempus partnership for benchmarking and research collaboration in Phase 3.
Build rationale	No vendor platform delivers a turnkey clinical outcomes intelligence engine for this patient population. The dataset is the strategic asset; the build is the mechanism for converting it into a product capability.
Estimated build timeline	6-9 months for Phase 2 internal analytics; Phase 3 external partner analytics built on the same foundation from month 12.

Databricks - Recommended Infrastructure Foundation

Market positioning: Healthcare-specific Lakehouse platform with strong BAA posture; the leading foundation for healthcare ML workloads as of Q3 2025.

Databricks is the recommended data infrastructure foundation for Use Cases 3, 4, and 5 simultaneously. The healthcare-specific Lakehouse architecture provides HIPAA-compliant storage and compute with BAA coverage, the MLflow integration layer supports the model training and deployment lifecycle for the predictive engagement model, and the SQL analytics layer supports the partner reporting automation. Selecting Databricks as the shared data platform across three use cases

eliminates the integration complexity of separate data environments per use case and reduces total infrastructure cost through platform consolidation. The alternative infrastructure candidate is Snowflake Healthcare Data Cloud, which offers comparable BAA posture and analytics capability with a different pricing model; the selection between Databricks and Snowflake should be made based on the engineering team's existing familiarity and the specific MLflow workflow requirements of the predictive modeling work.

Tempus - Strategic Research Partnership (Phase 3)

Market positioning: Healthcare AI data partnership platform with cardiometabolic research datasets and external benchmarking capabilities.

Tempus is recommended as a Phase 3 strategic partnership rather than a platform purchase. The Tempus data partnership provides cardiometabolic benchmarking data that complements the Client's own outcomes dataset, enabling the external-facing analytics capability to compare the Client's patient outcomes against a broader benchmark rather than solely the Client's internal historical data. This benchmarking layer is what makes the Phase 3 partner-facing analytics genuinely valuable to hospital partners: they want to know not just how their patients did, but how that compares to the broader population. The Tempus partnership should be initiated in Phase 2 to allow data integration and benchmarking calibration work to be completed before the Phase 3 external analytics launch.

> *The platforms named above reflect the healthcare AI market landscape as of Q3 2025. Specific platform selections require current-state evaluation during the implementation phase, including vendor financial stability, current product roadmap, integration capability with the Client's existing EHR and platform stack, and updated competitive positioning. The recommendation in this section is the build-buy-partner path and the evaluation criteria; the named platforms are the starting points for that evaluation.*

AI Infrastructure and Platform Recommendations

The five recommended use cases require a shared infrastructure foundation that does not exist in the Client's current technology environment. The infrastructure investment is not optional: without it, the recommended use cases cannot be built at production quality, at HIPAA-compliant scale, or with the observability required for clinical safety governance. The infrastructure investment is also not duplicative: the same Databricks data platform, the same BAA-governed LLM API layer, and the same observability tooling serve all five use cases. This section describes what that infrastructure is, why each component is necessary, and what it costs at the Client's scale.

Model Provider Strategy

The Client's AI applications require access to foundation models across two tiers. The first tier is commercial frontier LLM APIs with full HIPAA BAA coverage, used for the AI Care Companion conversational layer, the session intelligence note generation, the partner report narrative generation, and the outcomes intelligence engine's language generation components. The three primary candidates are OpenAI's Healthcare BAA offering (GPT-4o and GPT-4o-mini tiers), Anthropic's Enterprise agreement (Claude Sonnet and Claude Haiku tiers), and Google Cloud Healthcare API (Gemini Pro tier within the Google Cloud Healthcare data boundary). All three offer production HIPAA BAA coverage as of Q3 2025 with comparable capability levels. The selection among them should be based on pricing at the Client's inference volume, latency characteristics for the real-time conversational use cases, and the engineering team's existing familiarity.

The second tier is open-source models for cost-sensitive workloads, specifically the predictive engagement model and the clinical outcomes analytics layers that do not require general language generation capability. Open-source model deployment on the Client's cloud infrastructure eliminates per-token inference cost for these workloads and allows the engineering team to fine-tune on de-identified patient data without the data egress risk of commercial API calls. The recommended approach is to use commercial tier-one LLMs for user-facing conversational and generation workloads, and open-source models (Llama 3 class or equivalent) deployed on the Databricks ML infrastructure for analytical and predictive workloads.

A model router and policy engine governs which model class is used for which use case and enforces the data handling boundaries between PHI-eligible and de-identified workloads. This policy engine is a required component before any AI feature goes to production.

Inference Infrastructure and Cloud Deployment

The recommended cloud deployment is Microsoft Azure, specifically Azure Health Data Services, for the following reasons: Azure's BAA coverage for healthcare AI workloads is the most mature in the market as of Q3 2025, covering both the compute and data storage layers under a single BAA; the Databricks Lakehouse recommended for the data infrastructure layer has a native integration with Azure that simplifies the data flow architecture; and Azure's HIPAA compliance controls are well-documented and have been validated in peer healthcare AI deployments of comparable scale. AWS is a viable alternative if the engineering team has existing AWS infrastructure expertise; the BAA

posture is comparable. GCP is recommended only for organizations already operating within the Google Cloud Healthcare ecosystem.

Dedicated tenancy is recommended for the patient-facing AI Care Companion and session intelligence components. These components process PHI in real time and must be isolated from multi-tenant compute environments. The predictive modeling and outcomes analytics workloads can operate in standard HIPAA-eligible cloud environments without dedicated tenancy, reducing infrastructure cost for the non-real-time workloads.

Data Infrastructure Additions

Four specific data infrastructure components are required before the recommended use cases can be deployed to production. The Databricks Lakehouse is the foundational component, providing HIPAA-compliant storage, compute, and ML pipeline infrastructure for the outcomes dataset and the predictive modeling workloads. A vector store for clinical protocol retrieval is required for the AI Care Companion and session intelligence features; the recommended implementation is Databricks Vector Search, which operates within the same HIPAA boundary as the Lakehouse. A feature store for predictive modeling is required for the engagement churn model; Databricks Feature Store is the recommended implementation for the same platform consolidation reason. A secure de-identification pipeline for training data workloads is required before any patient data is used for model training; the assessment team recommends Microsoft Presidio for the initial implementation, with the de-identification pipeline reviewed by the clinical compliance team before training workloads begin.

Integration Patterns

The AI capabilities plug into the existing technology stack through four integration patterns. The EHR integration layer routes structured patient data (clinical measures, appointment history, care plan updates) from the existing EHR connections into the Databricks data environment through a standardized FHIR data extraction pipeline. The patient app integration layer connects the Flutter app to the AI Care Companion service through a RESTful API with end-to-end encryption and PHI access controls. The telehealth integration layer connects the telehealth platform to the AI session intelligence service through a real-time audio routing API, subject to the integration validation described in the technology readiness assessment. The partner reporting integration layer connects the Databricks reporting environment to the partner distribution workflow through a secure data output pipeline with partner-specific formatting configuration.

Observability and Evaluation Infrastructure

Model performance monitoring, evaluation harnesses, prompt management, and audit logging are not optional components in a clinical AI deployment. They are required by the governance framework described in Section 8 and by HIPAA's audit trail requirements. The recommended observability stack includes Databricks MLflow for model performance monitoring and experiment tracking, a custom evaluation harness for clinical safety review that samples AI output weekly for clinical team review, a prompt management system for the conversational AI layers that allows protocol updates to be deployed without redeployment cycles, and comprehensive audit logging that captures every

AI-generated output, every user interaction, and every clinical escalation event in a tamper-evident log. The audit log is the primary evidence base for clinical safety review and for payer audit defense.

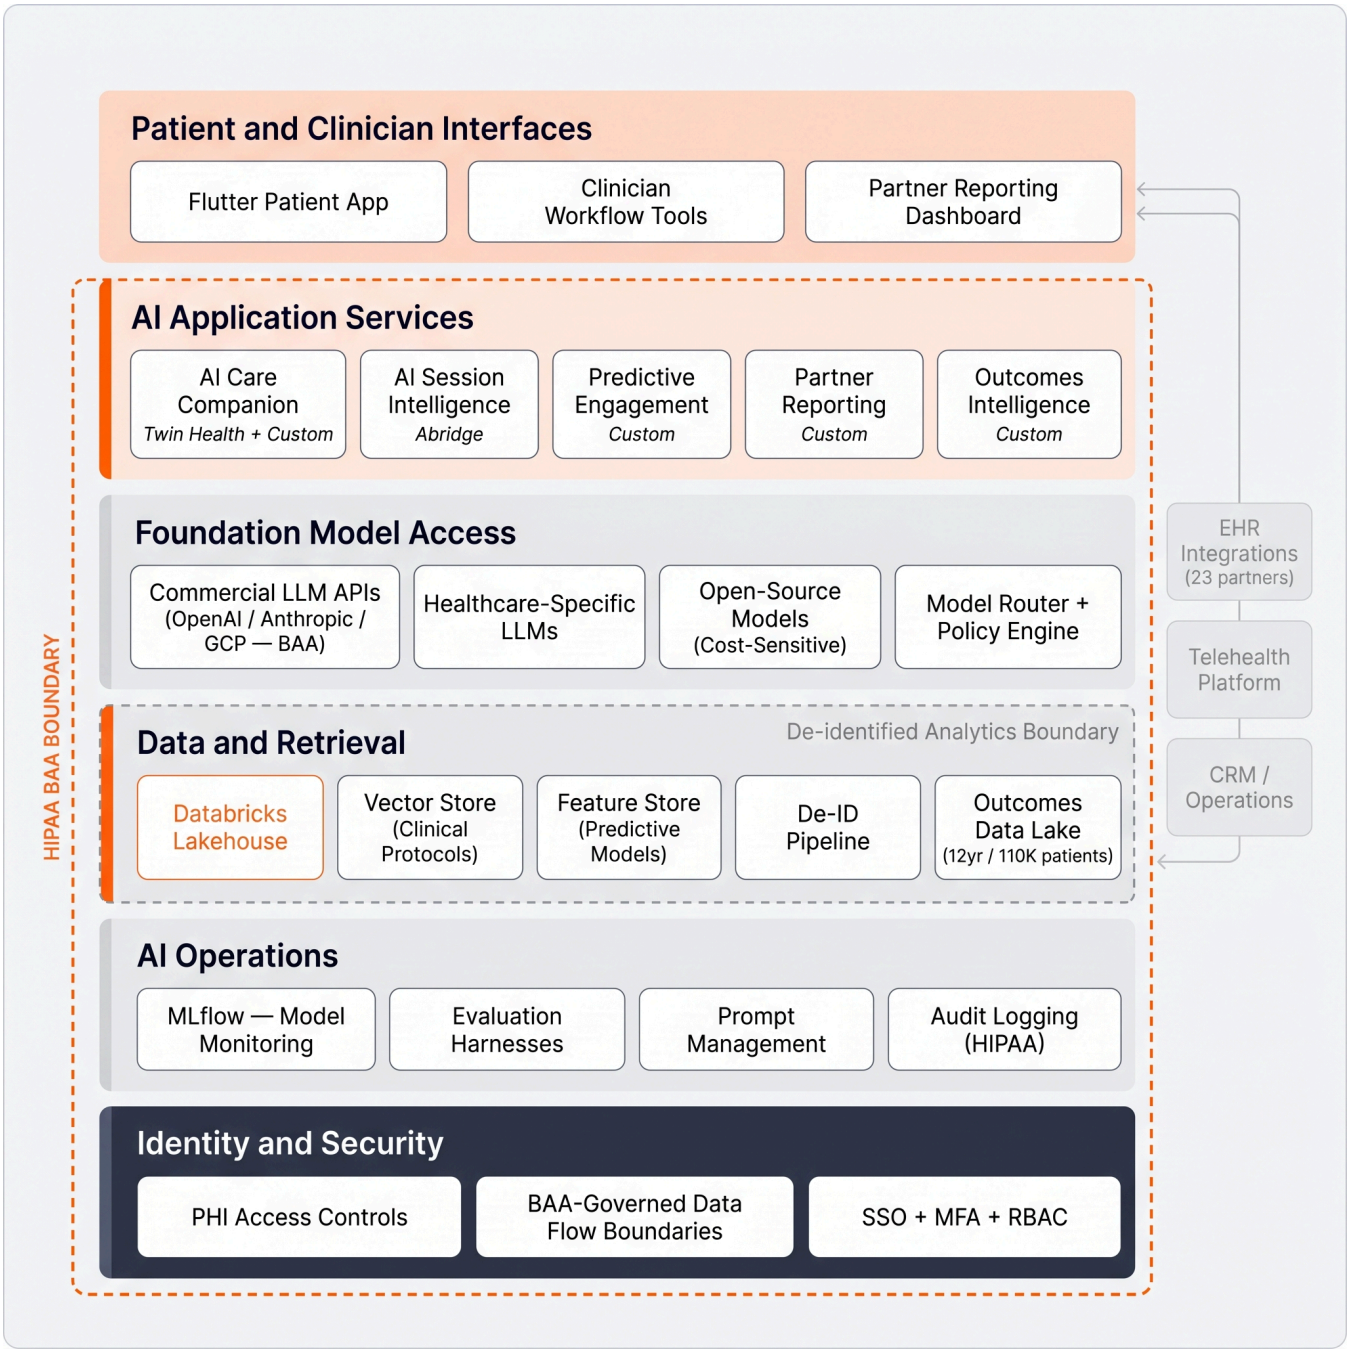

AI Governance and Risk Framework

The governance framework for the Client's AI program is proportionate to a mid-market clinical organization deploying AI in support of, not as a replacement for, licensed clinical judgment. It is not the governance framework of a hospital system operating FDA-cleared clinical decision support software. It is appropriate for an organization where every patient-facing AI feature has a licensed NP or RD in the loop, where no AI system makes autonomous clinical decisions, and where the primary regulatory exposure is HIPAA, payer audit, and state-specific clinical AI regulations that are still taking shape as of Q3 2025.

Model Risk Management Policy

No AI capability deploys to production without completing a structured review process. The review covers four dimensions: clinical safety (does the AI feature introduce any pathway by which a patient could receive incorrect or harmful clinical guidance without human review?), data quality (is the training data of sufficient quality and representativeness for the intended use?), bias and fairness (does the model perform consistently across the patient demographic subgroups present in the Client's patient population?), and regulatory alignment (does the feature description bring it within FDA SaMD jurisdiction or any state-specific clinical AI registration requirement?). Reviews are conducted by the AI Governance Committee described below and must be completed before any production deployment.

Clinical Safety and Patient Harm Prevention

The human-in-the-loop design principle is non-negotiable for all clinically-adjacent AI features. Specifically: the AI Care Companion does not provide diagnoses, does not modify care plans, and does not make autonomous decisions about clinical escalation; it surfaces escalation recommendations to the assigned NP or RD, who acts on them. The AI session intelligence does not write to the EHR autonomously; it produces a draft that the clinician reviews and approves. The predictive engagement model does not contact patients autonomously; it routes recommendations to the care coordination team, who make the outreach decision. The partner reporting automation does not send reports to partners autonomously; it produces report drafts that the clinical operations team reviews and approves.

Clinical output sampling is required for the AI Care Companion and session intelligence features. A random sample of 5% of AI Care Companion conversations and 10% of session intelligence note drafts are reviewed weekly by the clinical quality team during the first 6 months of production deployment. The review cadence is reduced to monthly after 6 months if the quality metrics remain within acceptable bounds. Any sample that reveals a clinically unsafe output triggers an immediate review by the AI Governance Committee and a temporary suspension of the affected feature pending remediation.

Data Privacy and PHI Protection

Patient data flows through the AI system architecture in two modes. PHI-eligible data flows through BAA-covered components only: the Azure Health Data Services environment, the commercial LLM

APIs with healthcare BAA coverage, and the Databricks Lakehouse with HIPAA-eligible configuration. De-identified data for model training and outcomes analytics flows through the secure de-identification pipeline before entering any analytical workload and is stored separately from the PHI-eligible environment. The de-identification methodology follows the HIPAA Safe Harbor standard; Expert Determination de-identification is used for the training dataset compiled from the 12-year outcomes data, given the small subpopulations in some geographic regions.

Audit logging captures all patient data access events, all AI-generated outputs that involve patient data, and all data export events. The audit log is reviewed quarterly by the compliance team and is available for payer audit on request. The consent management review recommended in the data readiness section is required before any patient data is used for the outcomes intelligence engine's analytical workloads.

Bias and Fairness Framework

The predictive engagement and churn model is the AI feature with the highest bias risk for this Client. If the model systematically underperforms for specific demographic subgroups, it will result in lower proactive outreach to patients in those subgroups, which translates to differential program completion rates by demographic. The assessment team reviewed the Client's historical drop-off data and identified two subgroups with historically higher drop-off rates: patients over 70 in the Hypertension Program and patients with lower health literacy scores in the Metabolic Program. The predictive model must be evaluated for subgroup performance on these dimensions before production deployment, and the model performance metrics must be disaggregated by these subgroup categories in all ongoing monitoring.

The AI Care Companion conversational layer carries a secondary bias risk: if the clinical protocols that anchor the AI responses embed historical clinical bias (for example, through under-representation of specific demographic groups in the historical outcomes data used to develop the protocols), the AI responses will reflect that bias. The clinical team's protocol review process must include an explicit bias audit before the protocols are converted to machine-readable format for the AI Care Companion.

Vendor and Model Provider Risk

The due diligence framework for healthcare AI vendors covers five dimensions: vendor financial stability (minimum 24 months of runway or profitable operations), BAA scope and indemnification terms (BAA must cover both data at rest and data in transit, with the vendor bearing breach indemnification for data within their custody), integration portability (can the Client's data be exported from the vendor platform in standard formats without significant friction?), product roadmap alignment (does the vendor's roadmap support the Client's telehealth and NP/RD encounter requirements?), and reference checks with comparable healthcare organizations. Vendor evaluation for Abridge and Twin Health specifically must include reference checks with virtual care organizations, not only inpatient or ambulatory health system deployments.

Regulatory Readiness

The AI features recommended in this strategy are designed to function as clinical decision support tools that augment licensed clinicians rather than as autonomous clinical decision makers. Under the current FDA AI/ML Software as a Medical Device framework, the human-in-the-loop design of all recommended features places them outside the FDA's current enforcement priorities for SaMD regulation, based on the assessment team's reading of the FDA's October 2024 AI/ML action plan and HHS Office for Civil Rights HIPAA AI guidance 2025. This assessment is not a legal opinion and should be confirmed by the Client's legal counsel before production deployment of the AI Care Companion and session intelligence features.

State-specific clinical AI regulations are an active area of legislative development as of Q3 2025. Several states have introduced or passed legislation requiring disclosure of AI use in clinical settings, transparency in algorithmic decision-making in healthcare, and consent for AI-assisted clinical interactions. The Client's 32-state operating footprint creates meaningful compliance surface area. The recommended approach is to implement a universal disclosure standard (patient-facing disclosure that AI tools support the clinical team's work) that satisfies the most restrictive current state requirements, rather than implementing a state-by-state compliance approach that becomes unmanageable at scale.

Governance Structure

The AI Governance Committee is the organizational mechanism for implementing this framework. The Committee meets monthly during Phases 1 and 2 and quarterly from Phase 3 onward, or at any time when a clinical safety trigger is activated. Membership: Chief Medical Officer (Chair), Head of Engineering, Head of Data Science, Chief Operating Officer, one NP representative elected by the clinical team, one member of the legal and compliance function, and one board advisor (the board member with healthcare operating experience). The Committee approves all new AI feature deployments, reviews clinical safety sampling reports, reviews model performance monitoring reports, and signs off on vendor selections before contract execution.

Risk Type	Mitigation Approach	Governance Owner	Review Cadence
Patient harm from AI Care Companion error	Human-in-the-loop escalation only; 5% weekly output sampling; immediate feature suspension trigger	CMO / AI Governance Committee	Weekly (sampling), Monthly (committee)
PHI exposure through LLM API	BAA-covered API only; no PHI in prompts to non-BAA models; audit logging of all API calls	Head of Engineering / Compliance	Quarterly audit
Bias in predictive engagement model	Subgroup performance evaluation pre-deployment; disaggregated monitoring metrics; quarterly bias review	Head of Data Science / CMO	Quarterly
Regulatory change (state AI laws)	Universal disclosure standard; quarterly regulatory landscape review; legal counsel briefing	CMO / Legal	Quarterly

Risk Type	Mitigation Approach	Governance Owner	Review Cadence
Vendor BAA breach	Vendor financial due diligence; BAA indemnification terms; portability requirements; rapid response protocol	Head of Engineering / Legal	As needed
Model performance degradation	MLflow monitoring; automated alerting at defined performance thresholds; monthly model performance review	Head of Data Science	Monthly
Clinical team protocol deviation from AI suggestion	Protocol adherence tracking; monthly clinical quality review; NP/RD feedback integration	CMO / Head of Clinical Operations	Monthly

Talent and Organizational Model

The AI strategy cannot be executed by the current engineering and data science teams without selective additions. The assessment team's evaluation of skills required versus skills present reveals three specific gaps: production AI engineering (the ability to build, deploy, and maintain ML models in production healthcare environments), clinical AI design (the ability to translate clinical protocols and clinical safety requirements into AI system design), and AI program management (the ability to manage vendor evaluation, cross-functional delivery, and governance coordination across a complex multi-use-case program). Each gap is addressable through a combination of targeted hiring and external partnership.

Required Skills by Use Case

Use Case	Engineering Skills Required	Current State	Gap
AI Care Companion	Flutter app development, LLM API integration, vector retrieval architecture, clinical protocol engineering	Flutter experience: 1 engineer. LLM integration: 0. Vector retrieval: 0.	Material - 2-3 engineers needed
AI Session Intelligence	Telehealth API integration, real-time audio processing, EHR write-back integration, vendor API	Telehealth integration: 2 engineers. Vendor API: requires partner onboarding.	Moderate - 1 engineer + vendor onboarding
Predictive Engagement	ML model development, Databricks, feature engineering on EHR data, workflow automation	EHR data: 1 data scientist. Databricks: 0. Production ML: 1 engineer.	Material - 1 ML engineer + Databricks expertise
Partner Reporting	Databricks, LLM integration for report generation, EHR FHIR extraction, report formatting	EHR integration: 2 engineers. LLM: 0. Databricks: 0.	Moderate - 1 engineer + LLM integration
Outcomes Intelligence Engine	Data science, ML engineering, Databricks ML, clinical data modeling	Data science: 3 analysts + 1 senior. Production ML: 1 engineer.	Moderate - 1 ML engineer + Databricks training

Recommended Organizational Model

The recommended structure is a small centralized AI team responsible for shared infrastructure, governance, and AI engineering standards, paired with embedded AI engineers in the product teams responsible for each use case. A Clinical AI Lead role bridges the engineering and clinical teams.

Role	Type	Primary Responsibility	Phase	Annual Cost (estimate)
Clinical AI Lead	Full-time hire	Translates clinical protocols to AI system design; represents clinical team requirements in engineering decisions; leads clinical safety sampling review	Phase 1 (Month 1)	$160K - $190K
Senior AI Engineer (Flutter + LLM)	Full-time hire	Leads Flutter app build and AI Care Companion engineering; LLM integration architecture	Phase 1 (Month 1)	$170K - $200K
ML Engineer (Databricks)	Full-time hire	Builds predictive engagement model and outcomes intelligence engine on Databricks; owns ML pipeline	Phase 1 (Month 2)	$165K - $195K
AI Program Manager	Full-time hire	Manages vendor evaluation, cross-functional delivery coordination, and governance committee operations	Phase 1 (Month 1)	$130K - $155K
AI Engineering Partner (Clixlogix)	External partnership	Staff augmentation for Phase 1 build acceleration; session intelligence integration; partner reporting build	Phase 1-2	$280K - $380K (Phase 1)
Data Engineer (EHR + FHIR)	Full-time hire or contract	FHIR extraction pipeline for all 23 partner EHR connections; de-identification pipeline build	Phase 1 (Month 2-3)	$140K - $165K

> *Note on cost benchmarks: the annual cost figures above reflect US market compensation ranges for comparable roles as of Q3 2025. They are presented at market-rate benchmarks to give the board a reference point for the full cost of building this capability with a domestic team. Clixlogix's India-based AI engineering practice delivers equivalent technical capability at substantially lower all-in cost, reflecting offshore delivery economics rather than US market salary structures. The AI Engineering Partner line reflects Clixlogix's engagement-based billing rather than full-time salary cost. Clients who wish to understand the Clixlogix-specific delivery cost model and how it compares to the US market benchmarks above are encouraged to request a detailed resource and cost plan during engagement scoping.*

Change Management for Clinical Teams

The NP and RD teams whose workflows will be most directly affected by AI session intelligence are the population whose adoption resistance represents the primary execution risk in Phase 1. The assessment team's interview data (NP and RD interviews, n=11, August 2025) identified three specific concerns that must be addressed in the change management approach: the concern that AI

documentation will reduce the quality of clinical notes relative to the clinician's own narrative; the concern that AI session suggestions will create pressure to follow protocol prompts even when clinical judgment would recommend deviation; and the concern that AI monitoring of sessions creates a surveillance dynamic that affects the clinician-patient relationship.

The change management approach addresses each concern specifically. Co-design participation: NPs and RDs from each program line participate in the design and testing of the session intelligence features before production deployment. Their specific note-format preferences and protocol-deviation documentation needs are incorporated into the feature design, not retrofitted after the fact. Opt-in phasing: Phase 1 session intelligence deployment begins with the volunteer cohort identified in interviews. Participation is not mandatory during Phase 1. Mandatory rollout, if required, follows Phase 1 evidence of quality improvement. Clinical sovereignty framing: the AI session intelligence is positioned in all clinical communications as a tool that gives the clinician time back, not as an oversight mechanism. Clinical leadership (CMO and Head of Clinical Operations) must be visible champions of this framing from the first communication.

Investment Envelope and ROI

The financial projections in this section are built from the bottom up by use case and infrastructure component, using the Client's actual unit economics and the assessment team's vendor pricing research at the Client's scale. The projections carry explicit uncertainty ranges where the assumptions are most sensitive, and the section closes with a sensitivity analysis that frames the conditions under which the financial case strengthens or weakens materially.

Capital Investment by Use Case (18-Month Program)

Use Case / Component	Phase 1 (M0-6)	Phase 2 (M6-12)	Phase 3 (M12-18)	Total
Flutter patient app build (platform prerequisite)	$320K - $380K	-	-	$320K - $380K
AI Care Companion (custom build + Twin Health pilot)	$180K - $220K	$280K - $340K	$80K - $100K	$540K - $660K
AI Session Intelligence (Abridge integration + build)	$120K - $150K	$80K - $100K	-	$200K - $250K
Predictive Engagement Model (Databricks + build)	$240K - $290K	$120K - $150K	$60K - $80K	$420K - $520K
Partner Reporting Automation (build)	$180K - $220K	$160K - $200K	$40K - $60K	$380K - $480K
Clinical Outcomes Intelligence Engine (Databricks + build)	-	$320K - $400K	$240K - $300K	$560K - $700K
Databricks Lakehouse setup + data infrastructure	$280K - $340K	$80K - $100K	$40K - $60K	$400K - $500K
External AI engineering partnership (Clixlogix)	$280K - $380K	$160K - $240K	-	$440K - $620K
Total Capital Investment	$1.6M - $1.98M	$1.2M - $1.53M	$0.46K - $0.6M	$3.26M - $4.11M

Operating Cost Projection (Annualized at Month 12)

Cost Category	Annual Cost (Month 12 run rate)	Notes
Abridge licensing (120 NP and RD users)	$360K - $480K/yr	Estimate based on mid-market ambient documentation pricing; subject to vendor negotiation

Cost Category	Annual Cost (Month 12 run rate)	Notes
Twin Health pilot licensing (Metabolic Program subset)	$120K - $180K/yr	Phase 1 pilot scale; full program licensing to be negotiated based on pilot outcomes
Databricks Lakehouse (compute + storage)	$180K - $240K/yr	At Client's data scale; Databricks Healthcare pricing Q3 2025
Commercial LLM API (OpenAI / Anthropic Healthcare BAA)	$80K - $140K/yr	Dependent on AI Care Companion usage volume; inference cost trajectory likely declining
Azure Health Data Services (infrastructure)	$60K - $90K/yr	Dedicated tenancy for patient-facing components; standard HIPAA for analytics
New FTE additions (4-5 hires by Month 12)	$620K - $750K/yr	Fully-loaded including benefits; based on market compensation ranges Q3 2025
Total Operating Cost (annualized at Month 12)	$1.42M - $1.88M/yr	Excludes Tempus partnership (Phase 3 cost); excludes revenue cycle or consumer AI

Expected Business Impact by Use Case

Use Case	Impact Type	12-Month Estimate	18-Month Estimate	Basis
AI Care Companion	Revenue (Metabolic + Hypertension retention)	$1.4M - $2.1M	$2.8M - $4.1M	12-15% completion rate improvement; Client per-patient economics
AI Session Intelligence	Cost (capacity unlock)	$0.8M - $1.2M	$1.6M - $2.4M	10-min/session documentation reduction; 80 NPs; Cardiac growth capacity
Predictive Engagement	Revenue (retention recovery)	$0.8M - $1.2M	$1.6M - $2.4M	40-50% churn event detection; 50% conversion to retention
Partner Reporting Automation	Cost + Revenue (operational + new partners)	$0.4M - $0.7M	$1.2M - $2.4M	1.5 FTE redeployed + 2-3 new partner

Use Case	Impact Type	12-Month Estimate	18-Month Estimate	Basis
				relationships in 12 months
Clinical Outcomes Intelligence Engine	Revenue (program design + partner analytics)	-	$0.6M - $1.2M	2-4% completion rate improvement; Phase 3 partner analytics positioning
Total Expected Impact		$3.4M - $5.2M	$7.8M - $12.5M	Ranges reflect assumption sensitivity; see sensitivity analysis below

Payback Period and ROI

At the midpoint of the capital and operating cost ranges and the midpoint of the expected impact ranges, the 18-month program generates approximately $10.2M in cumulative business impact against a total 18-month program cost of approximately $6.6M, producing a net benefit of approximately $3.6M over the program period. Payback is estimated at approximately 14-16 months from program authorization, with Phase 1 use cases (session intelligence and predictive engagement) generating the earliest return.

The three-year economic projection (extending the Phase 3 capabilities through months 18-36) projects cumulative impact in the $18M-$28M range against cumulative program cost of approximately $10M-$13M, including the Phase 3 extensions. The wide range on the three-year projection reflects the uncertainty in the partner analytics revenue contribution, which is the highest-expected-value but least-certain component of the long-term impact.

Sensitivity Analysis

Scenario	Variable Changed	Impact on Net Benefit (18-Month)	Assessment
Model inference costs fall 40% by Month 12	LLM API cost reduction (consistent with 2024-2025 market trajectory)	+$48K - $56K annually	Modest positive impact; inference is not the dominant cost
AI Care Companion completion rate improvement is 8% not 12%	Lower end of engagement improvement estimate	-$700K - $1.1M at 18 months	Manageable; Phase 1 predictive engagement and session intelligence ROI absorbs shortfall

Scenario	Variable Changed	Impact on Net Benefit (18-Month)	Assessment
Partner reporting adds 1 new partner (not 2-3)	Lower end of new partner acquisition estimate	-$1.2M - $1.8M at 18 months	Meaningful sensitivity; partner acquisition is the most uncertain impact driver
Abridge fails NP/RD telehealth evaluation; build path required	Session intelligence converts from buy to build	-$150K capital + 3-4 month delay	Execution risk; addressed by parallel build scoping in Phase 1
Competitive acceleration forces faster Metabolic re-entry	Front-load Twin Health pilot and Flutter app	+$80K - $120K Phase 1 cost	Manageable; trade-off is Phase 1 cost increase for earlier Metabolic Program defense

Where the Financial Projections Are Most Uncertain

Two impact drivers carry the highest uncertainty and warrant explicit board acknowledgment. First, the partner acquisition impact: the projection that automated partner reporting enables 2-3 new hospital partner relationships within 12 months is based on business development interview data indicating that reporting quality is a named friction point for prospective partners. The conversion from reduced friction to closed partnership depends on commercial execution factors beyond the AI investment itself. The assessment team recommends treating the partner acquisition impact as an upside scenario rather than a base case for financial planning purposes. Second, the Hypertension Program rescue: the predictive engagement model's effectiveness in the Hypertension Program is more uncertain than in the Metabolic Program because the Hypertension Program data quality gaps identified in the readiness assessment will limit model performance during Phase 1. The Phase 1 Hypertension engagement targets are accordingly conservative, and the Phase 2 extension depends on Phase 1 data quality remediation proceeding as planned.

Phased 18-Month Roadmap

Phase 1: Foundation and Proof (Months 0-6)

Phase 1 establishes the organizational, technical, and governance foundations without which nothing in Phase 2 is possible, and delivers the two use cases with the fastest time-to-value: predictive patient engagement for the Metabolic Program and AI session intelligence in a clinical pilot. Phase 1 also begins the Flutter app build that is the prerequisite for the AI Care Companion, and establishes the partner reporting automation MVP for the four largest hospital partner relationships.

Workstream	Phase 1 Deliverables (M0-6)	Key Milestone	Success Criteria
Platform Investment	Flutter app build initiated; Databricks Lakehouse provisioned on Azure; de-identification pipeline built and clinically reviewed; FHIR extraction from all 23 partner EHRs scoped and initiated for the 4 largest	Month 3: Flutter alpha; Month 5: Databricks live with historical data ingested	Flutter app in beta testing with 50 internal users by Month 5; Databricks ingesting 3+ years of outcomes data without quality loss
Patient-Facing AI	Twin Health evaluation and pilot contract executed; pilot deployment to 200-300 Metabolic Program patients; patient experience data collected	Month 4: Twin Health pilot live; Month 6: pilot evaluation report	Twin Health pilot achieving 40%+ patient engagement with the platform among enrolled patients; protocol extension feasibility confirmed
Clinician-Facing AI	Abridge and Ambience evaluated against NP/RD telehealth encounter criteria; winning vendor contracted; integration with telehealth platform built; pilot deployment to 10-15 volunteer NPs across Cardiac and Metabolic programs	Month 2: Vendor evaluation complete; Month 4: Pilot live with volunteer NP cohort	Post-session documentation time reduced by 40%+ among pilot NPs; no adverse clinical quality findings in output sampling
Operational AI	Predictive engagement model trained on Metabolic Program data; model deployed to Metabolic Program care coordination workflow; partner reporting MVP built for 4 largest partners	Month 4: Predictive model in production for Metabolic Program; Month 6: Partner reporting MVP live	60-day Metabolic Program retention rate tracking improvement; partner reporting delivery time reduced from 8-12 days to 2-3 days for 4 pilot partners
Governance	AI Governance Committee established and first meeting held; clinical safety sampling framework implemented; vendor due diligence framework documented; NP/RD change	Month 1: Committee chartered; Month 2: First full meeting	Committee meeting monthly; first vendor selection completed with documented due diligence; NP change management

Workstream	Phase 1 Deliverables (M0-6)	Key Milestone	Success Criteria
	management communications launched		sentiment survey showing 60%+ positive or neutral

Phase 2: Scale and Rescue (Months 6-12)

Phase 2 deploys the AI Care Companion to patients through the completed Flutter app, extends session intelligence to the full clinical team, scales the predictive engagement model to the Hypertension Program, and builds the first version of the clinical outcomes intelligence engine. Phase 2 is where the Hypertension Program rescue begins in earnest and where the Metabolic Program re-acceleration transitions from pilot to full deployment.

Workstream	Phase 2 Deliverables (M6-12)	Key Milestone	Success Criteria
Patient-Facing AI	AI Care Companion launched in Flutter app for Metabolic Program; Cardiac Program Care Companion rollout begins; Twin Health pilot findings incorporated into custom care companion design	Month 8: Care Companion live for Metabolic Program; Month 11: Cardiac Program rollout begins	15%+ improvement in 90-day Metabolic Program completion rate vs. Phase 1 baseline; Care Companion session engagement rate 55%+ of enrolled patients
Clinician-Facing AI	Session intelligence extended to all 80 NPs and 40 RDs; compliance reporting integrated into partner reporting workflow; full clinical quality monitoring operational	Month 8: Full clinical team rollout complete	Post-session documentation time reduced 40%+ across full clinical team; compliance reporting generated automatically for all sessions
Operational AI	Predictive engagement model extended to Hypertension Program; automated outreach workflow live for both Metabolic and Hypertension programs; partner reporting extended to all 23 partners	Month 9: Hypertension predictive model live; Month 10: All 23 partners on automated reporting	Hypertension 60-day retention rate improvement of 8%+ vs. Phase 1 baseline; all 23 partners receiving automated reports within 3 business days of month close
Distinctive AI	Clinical Outcomes Intelligence Engine Phase 2 build on Databricks; first internal program design insights delivered to CMO and clinical leadership; Tempus partnership initiated	Month 11: First outcomes intelligence dashboard live for internal use	Clinical team reporting active use of outcomes insights in protocol review sessions; Tempus data partnership agreement executed

Workstream	Phase 2 Deliverables (M6-12)	Key Milestone	Success Criteria
Governance	AI Governance Committee reviews Phase 2 deployments; bias review for predictive engagement model at full scale; regulatory landscape review updated for state AI law developments	Month 12: Phase 2 governance review and board update	All Phase 2 deployments reviewed and approved by Committee; no open clinical safety findings; Phase 3 authorized by board based on Phase 2 milestones

Phase 3: Distinctive Capability and Expansion (Months 12-18)

Phase 3 converts the internal capabilities built in Phases 1 and 2 into external-facing assets that differentiate the Client's partner value proposition and establish the strategic moat described in the AI thesis. Phase 3 also extends the AI Care Companion to the Hypertension Program and initiates the planning for months 18-36.

Workstream	Phase 3 Deliverables (M12-18)	Key Milestone	Success Criteria
Patient-Facing AI	AI Care Companion launched for Hypertension Program with adapted clinical protocols; protocol adaptations informed by Phase 2 outcomes intelligence findings	Month 15: Hypertension Care Companion live	Hypertension 90-day completion rate improving by Month 18; Care Companion engagement rate comparable to Metabolic Program performance
Clinician-Facing AI	Session intelligence in full steady-state operation; documentation efficiency gains documented for board reporting; next-generation session intelligence features scoped	Month 18: Full program documentation efficiency report	NP capacity buffer quantified and documented; Cardiac Program growth absorbed without emergency hiring
Operational AI	Partner reporting in steady-state for all 23 partners; partner satisfaction data collected; new partner onboarding workflow incorporating automated reporting from day one	Month 16: New partner onboarding process live with automated reporting	2+ new hospital partner relationships closed in which automated reporting was cited as a decision factor
Distinctive AI	Clinical Outcomes Intelligence Engine extended to external partner-facing analytics; first external outcomes intelligence reports delivered to hospital partners; Tempus benchmarking data integrated; 18-36 month	Month 17: First external outcomes intelligence reports delivered to pilot hospital partners	3+ hospital partners actively using external outcomes intelligence reports; partner renewal conversations citing

Workstream	Phase 3 Deliverables (M12-18)	Key Milestone	Success Criteria
	outcomes product roadmap drafted		outcomes analytics as a retention factor
Planning	18-36 month AI roadmap presented to board; fourth program line AI strategy optionality assessed; revenue cycle AI workstream scoped	Month 18: Board presentation of next roadmap phase	Board approval of 18-36 month roadmap; capital envelope for next phase authorized

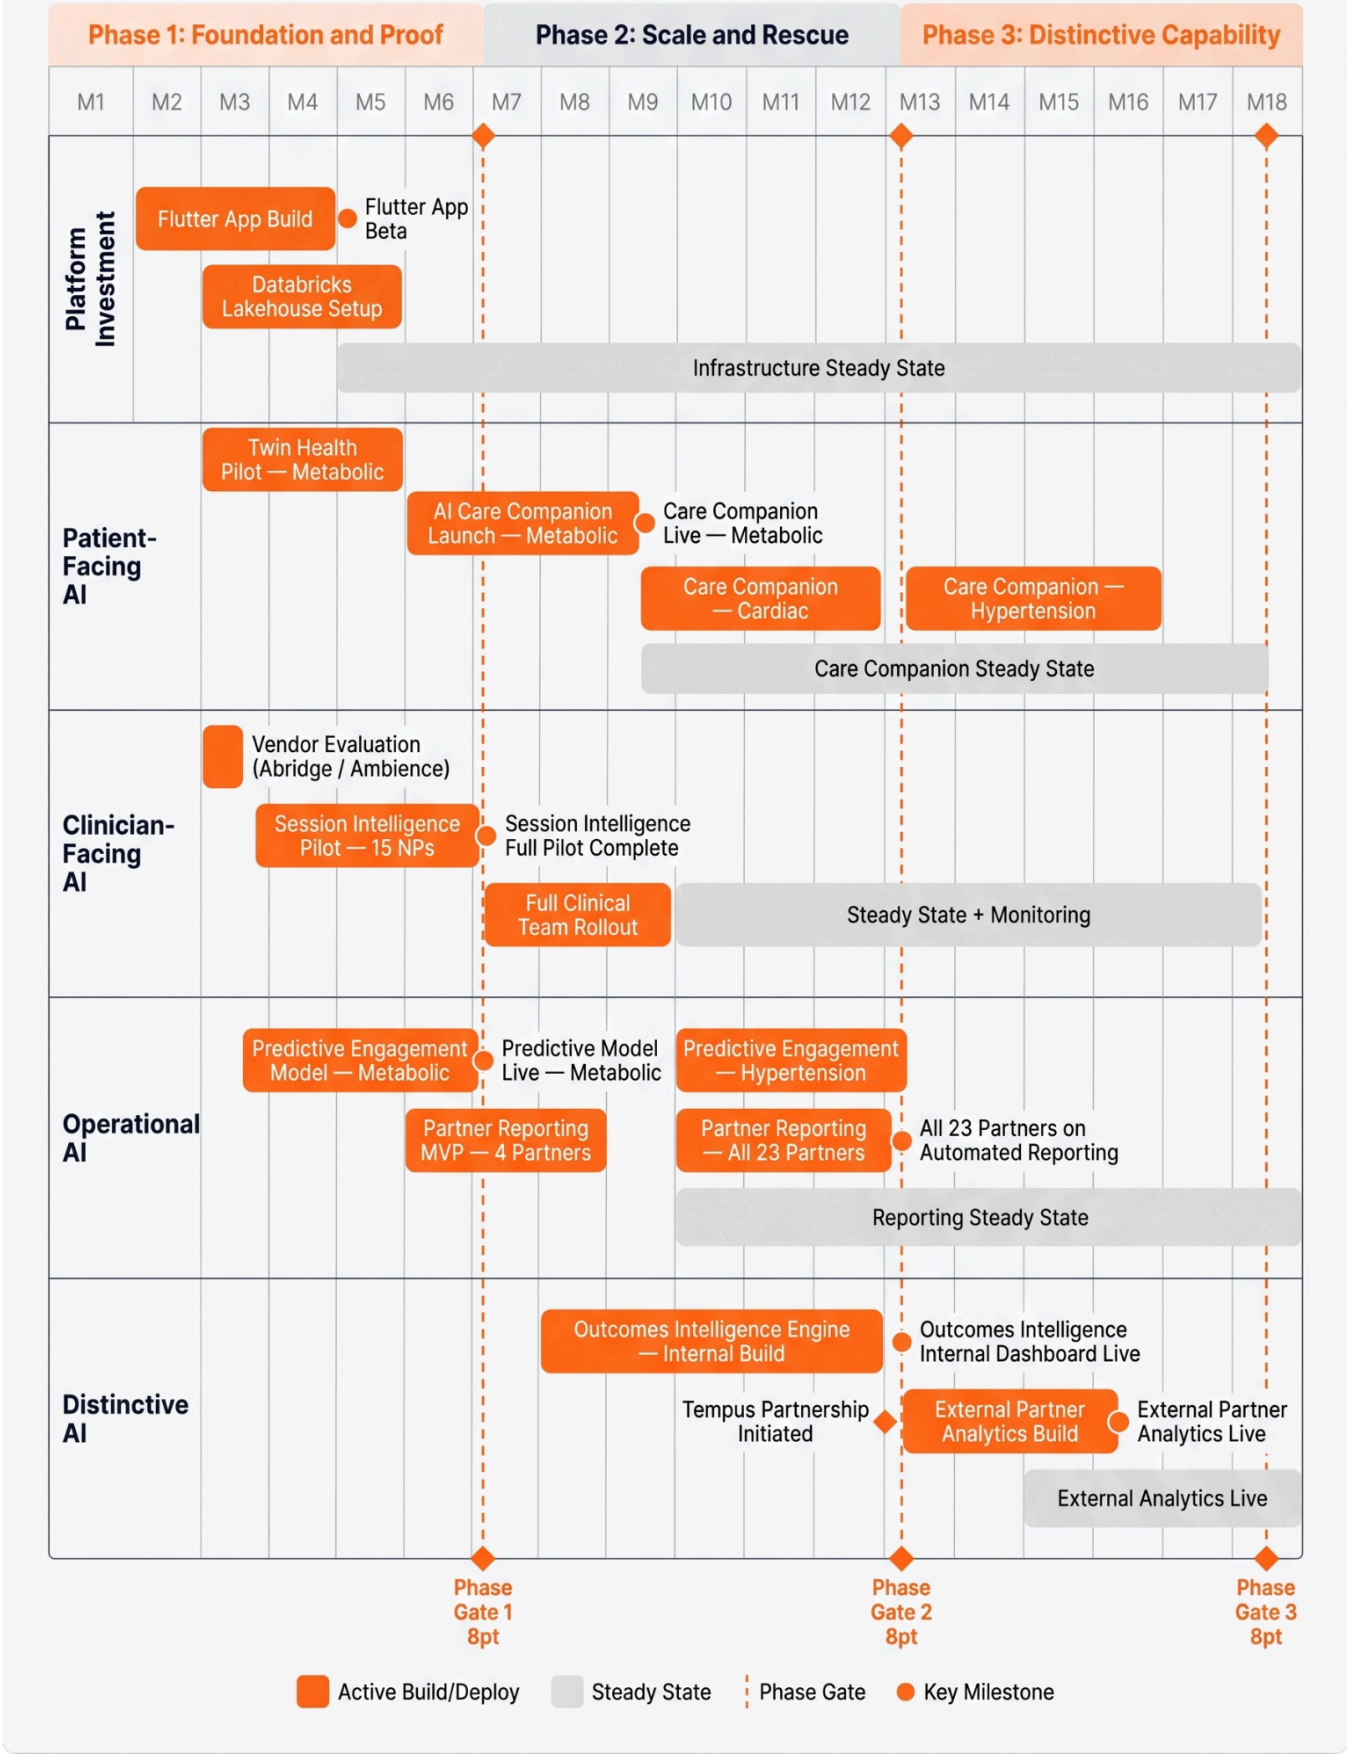
Phase 1: Foundation and Proof
Phase 2: Scale and Rescue
Phase 3: Distinctive Capability
M1
M2
M3
M4
M5
M6
M7
M8
M9
M10
M11
M12
M13
M14
M15
M16
M17
M18
Platform Investment
Flutter App Build
Flutter App Beta
Databricks Lakehouse Setup
Infrastructure Steady State
Patient-Facing AI
Twin Health Pilot — Metabolic
AI Care Companion Launch — Metabolic
Care Companion Live — Metabolic
Care Companion — Cardiac
Care Companion — Hypertension
Care Companion Steady State
Clinician-Facing AI
Vendor Evaluation (Abridge / Ambience)
Session Intelligence Pilot — 15 NPs
Session Intelligence Full Pilot Complete
Full Clinical Team Rollout
Steady State + Monitoring
Operational AI
Predictive Engagement Model — Metabolic
Predictive Model Live — Metabolic
Predictive Engagement — Hypertension
Partner Reporting MVP — 4 Partners
Partner Reporting — All 23 Partners
All 23 Partners on Automated Reporting
Reporting Steady State
Distinctive AI
Outcomes Intelligence Engine — Internal Build
Outcomes Intelligence Internal Dashboard Live
Tempus Partnership Initiated
External Partner Analytics Build
External Partner Analytics Live
External Analytics Live
Phase Gate 1 8pt
Phase Gate 2 8pt
Phase Gate 3 8pt
Active Build/Deploy
Steady State
Phase Gate
Key Milestone

Risk Register

The following risks are rated on a 1-5 scale for both severity (the impact on program outcomes if the risk materializes) and likelihood (the probability of materialization given current information). Risk ratings are as of Q3 2025; they should be reviewed at each phase gate.

#	Risk	Type	Severity	Likelihood	Risk Score	Owner	Mitigation
1	NP and RD adoption resistance to AI session intelligence delays Phase 1 pilot timeline	Execution	4	3	12 / High	CMO / Clinical Ops	Clinical co-design participation; opt-in phasing; clinical sovereignty framing; visible CMO sponsorship
2	Abridge and Ambience fail NP/RD telehealth encounter evaluation; build path required	Technical	3	3	9 / Med	Head of Engineering	Parallel build scoping begins in Month 1; telehealth encounter validation is Month 2 gate criterion
3	Flutter app build runs 2+ months over schedule; AI Care Companion deployment delayed	Technical	4	2	8 / Med	Head of Engineering	Dedicated senior Flutter engineer hired Month 1; weekly build progress review; MVP-first scoping
4	Twin Health pilot reveals protocol extension is not feasible; custom build required in Phase 1	Vendor Fit	3	2	6 / Med	Clinical AI Lead	Protocol extension feasibility is Month 2 evaluation gate; custom build design runs in parallel from Month 1
5	Databricks Lakehouse data ingestion reveals Hypertension Program data quality worse than assessed	Data	3	3	9 / Med	Head of Data Science	Hypertension data quality remediation scoped in Phase 1 alongside Databricks

#	Risk	Type	Severity	Likelihood	Risk Score	Owner	Mitigation
							build; model deployment gated on quality threshold
6	Commercial LLM API inference costs exceed projection by 50%+ at production scale	Technical / Financial	2	2	4 / Low	Head of Engineering	Open-source model evaluation for cost-sensitive workloads; inference cost monitoring from first deployment
7	State AI law enacted in key operating state requiring consent re-architecture	Regulatory	3	2	6 / Med	CMO / Legal	Universal disclosure standard implemented from Phase 1; quarterly regulatory landscape review; legal counsel on retainer
8	Competitive platform launches AI Care Companion equivalent in Metabolic market before Phase 2	Strategic	3	3	9 / Med	CEO / Business Development	Flutter app build and Twin Health pilot front-loaded in Phase 1; proprietary outcomes data moat not replicable in 6 months
9	Hospital partner pushback on AI-assisted reporting (concerns about data use or AI accuracy)	Strategic	3	2	6 / Med	VP Business Development / CMO	Partner communication strategy developed in Phase 1; pilot partners selected for AI affinity; partner-facing transparency on AI use
10	Key hire (Clinical AI Lead or Senior AI Engineer) takes 4+ months to fill	Execution	3	3	9 / Med	CEO / COO	Recruiting launched immediately at program authorization; external AI engineering

#	Risk	Type	Severity	Likelihood	Risk Score	Owner	Mitigation
							partnership (Clixlogix) bridges gap during search
11	Abridge or Twin Health vendor financial instability or acquisition during program	Vendor	3	2	6 / Med	AI Governance Committee	Vendor financial due diligence at selection; BAA portability requirements in contract; data export provisions required
12	FDA expands SaMD enforcement to cover AI Care Companion features	Regulatory	4	1	4 / Low	CMO / Legal	Human-in-the-loop design provides current regulatory buffer; legal counsel monitors FDA AI/ML action plan updates quarterly
13	Patient data breach through BAA-governed LLM API	Security	5	1	5 / Med	Head of Engineering / Legal	BAA-covered APIs only; PHI audit logging; incident response protocol; cyber insurance review
14	Predictive engagement model bias creates differential outreach rates across demographic subgroups	Clinical / Ethical	4	2	8 / Med	Head of Data Science / CMO	Pre-deployment subgroup performance evaluation; disaggregated monitoring metrics; quarterly bias review by AI Governance Committee

Recommendation Summary

The AI Thesis Restated

The Client's three durable assets -- 12 years of clinical outcomes data, a multidisciplinary clinical team embedded across 32 states, and 23 established hospital partner relationships -- are sufficient to build a defensible AI position if the investment is made now and executed with discipline. The AI thesis is not about becoming an AI company. It is about using AI to extend clinical capacity, reduce between-appointment patient drop-off, convert the partner reporting liability into a competitive asset, and build an outcomes intelligence capability that no competitor can replicate without the same 12 years of data. The five recommended use cases execute that thesis across three programs and one partner relationship in a phased, evidence-gated sequence.

The Do-Nothing Alternative

Not investing in AI is not a neutral decision. Two well-funded competitors have already deployed AI-powered patient features in the past six months. The Metabolic Program is flat year-over-year under competitive pressure from platforms with AI health companions and AI-assisted clinical notes. The Hypertension Program is contracting. The clinical team's documentation burden will grow linearly with the Cardiac Program's 40% annual growth unless AI absorbs it. The partner reporting process will continue to consume 1.5 FTE of clinical operations capacity and limit new partner onboarding. The 12-year outcomes dataset will continue to appreciate in retrospective value while generating zero current competitive advantage. In 18 months, the competitive gap will be materially harder to close.

The Do-Everything Alternative

Deploying all 20 candidate use cases simultaneously would consume more capital than the AI strategy warrants, more engineering capacity than the team has, more change management bandwidth than the clinical team can absorb, and more governance attention than a mid-market clinical organization can provide without clinical safety risk. The recommended five use cases were selected precisely because they are the highest-impact, most feasible, most strategically differentiated set that can be executed within organizational capacity. Adding use cases to Phase 1 without adding capacity creates delivery risk that propagates to all five recommended use cases.

Conditions for Revisiting the Recommendation

The board should revisit this recommendation under three specific conditions. First, if Phase 1 milestones at month 6 are not met -- specifically, if the predictive engagement model is not in production for the Metabolic Program, the session intelligence pilot has not demonstrated 40%+ documentation time reduction, and the Flutter app is not in beta -- the Phase 2 authorization should be conditional on a revised plan rather than automatic. Second, if a competitive development materially changes the threat landscape before Phase 2 authorization (for example, if a competitor acquires a clinical outcomes dataset of comparable scale), the build-versus-buy stance for the outcomes intelligence engine should be re-evaluated. Third, if the FDA issues specific guidance on AI-assisted

clinical programs that brings any of the recommended features within SaMD jurisdiction, the affected features should be paused pending legal review before Phase 2 deployment.

Appendix

A. Use Case Scoring Detail

The following table provides the complete scoring detail for the top eight use cases evaluated in the prioritization framework. Scores are on a 1-5 scale per dimension; the weighted score applies the framework weights (35% impact, 30% feasibility, 25% strategic fit, 10% time-to-value).

Use Case	Impact (35%)	Feasibility (30%)	Strategic Fit (25%)	Time-to-Value (10%)	Weighted Score
AI Session Intelligence	4.5	4.0	5.0	4.0	4.3
Predictive Patient Engagement	4.5	4.0	4.5	4.0	4.3
AI Care Companion	5.0	3.5	5.0	3.0	4.1
Partner Reporting Automation	4.0	3.5	4.5	3.5	3.9
Clinical Outcomes Intelligence	5.0	3.0	5.0	2.0	3.8
AI-Powered Intake and Onboarding	3.0	3.0	3.0	3.5	3.0
NP Panel Optimization Engine	3.0	2.5	3.0	2.5	2.8
Hypertension Protocol AI	3.0	2.5	3.5	2.5	2.9

B. Vendor Evaluation Summary

The following table summarizes the primary evaluation criteria applied to the healthcare AI vendors assessed during this engagement. Criteria weightings reflect this Client's specific priorities as a virtual specialty care organization with NP and RD clinical staff operating in telehealth settings.

Vendor	Category	HIPAA BAA	NP/RD Telehealth	Mid-Market Fit	Data Portability	Financial Stability	Recommended Status
Abridge	Ambient AI Documentation	Confirmed	Requires validation	Strong	Standard export	Strong - Best in KLAS Ambient AI 2025 and 2026; $550M raised; 200+ health system deployments	Primary evaluation

Vendor	Category	HIPAA BAA	NP/RD Telehealth	Mid-Market Fit	Data Portability	Financial Stability	Recommended Status
Ambience	Ambient AI Documentation + CDI + Coding	Confirmed	Requires validation	Strong - Cleveland Clinic win	Standard export	Unicorn $1.25B valuation (Series C, July 2025); $345M raised	Alternate evaluation
Twin Health	Patient Engagement AI - Metabolic	Confirmed	N/A (patient-facing)	Strong for Metabolic / employer	Requires negotiation	$53M investment (Aug 2025); NEJM Catalyst RCT published; outcomes-based pricing model	Phase 1 pilot
Innovaccer (incl. Story Health)	Healthcare Intelligence Cloud: data (Gravity), analytics (Atlas), engagement (Cured), specialty care (Story Health by Innovaccer)	Confirmed	N/A	Mid-to-large market; serves 6 of top 10 US health systems	Standard FHIR export	Strong - $275M Series F (Jan 2025); Kaiser Permanente and Banner Health as strategic investors; 2026 Best in KLAS for Providers, Payers, and CRM	Conditional alternative - platform-led path
Health Catalyst	Population Health Analytics	Confirmed	N/A	Enterprise-oriented	Standard export	Publicly traded	Conditional alternative
Arcadia	Population Health Reporting	Confirmed	N/A	Mid-market strong	Standard export	Growth stage	Conditional alternative
Databricks	Data + ML Infrastructure	Confirmed	N/A	All market segments	Open format (Delta Lake)	Publicly traded	Recommended infrastructure
Tempus	Data Partnership - Cardiometabolic Benchmarking	Confirmed	N/A	Research partnership	Data collaboration agreement	Growth stage	Phase 3 partnership

C. Regulatory and Framework References

The following regulatory frameworks and research references informed the governance recommendations and regulatory readiness assessment in this strategy:

- HHS Office for Civil Rights, HIPAA AI Guidance, 2025. Guidance on the application of HIPAA Privacy Rule requirements to AI-assisted healthcare functions.
- FDA Center for Devices and Radiological Health, AI/ML-Based Software as a Medical Device Action Plan, Updated 2024. Framework for determining SaMD jurisdiction for AI clinical decision support tools.
- HIMSS 2025 State of Healthcare AI Survey. Industry benchmark data on AI adoption rates, implementation challenges, and organizational readiness in US healthcare organizations.
- Rock Health 2025 Digital Health Funding Report. Funding landscape for digital health and healthcare AI vendors; referenced for vendor financial stability context.
- Menlo Ventures 2025 State of AI in Healthcare Report. Healthcare AI use case adoption data and patient engagement benchmarks referenced in impact projection methodology.
- KLAS Research 2025 Best in KLAS Awards. Ambient clinical documentation category rankings referenced in Section 6 vendor assessment.
- JAMA 2024 AI in Clinical Practice Review. Peer-reviewed evidence base for human-in-the-loop clinical AI design principles referenced in governance framework.
- Clixlogix Consulting Practice, Healthcare AI Readiness Assessment Framework (v2.4), 2025. Proprietary five-dimension readiness assessment framework applied in Section 4.

D. Glossary

Term	Definition
Ambient Clinical Documentation	AI-assisted transcription and note generation during clinical encounters, producing structured clinical note drafts from session audio without requiring the clinician to type or dictate.
BAA (Business Associate Agreement)	A contract required under HIPAA between a covered entity (the Client) and a business associate (a vendor) that handles protected health information on behalf of the covered entity. BAA coverage is required for all vendor platforms that process patient data.
FHIR (Fast Healthcare Interoperability Resources)	A standard for healthcare data exchange developed by HL7. FHIR APIs are the current standard for interoperability between EHR systems and external applications.
Human-in-the-Loop	An AI system design in which a human reviews and approves AI-generated outputs before they take effect in the workflow. All AI features recommended in this strategy follow human-in-the-loop design.
Lakehouse	A data architecture that combines the structured query capabilities of a data warehouse with the flexible storage and ML capabilities of a data lake. Databricks Lakehouse is the specific implementation recommended in this strategy.
LLM (Large Language Model)	A foundation AI model trained on large text datasets, capable of generating, summarizing, and classifying natural language. Commercial

Term	Definition
	LLMs referenced in this strategy (GPT-4o, Claude, Gemini) are accessed via API under HIPAA-compliant BAA arrangements.
PHI (Protected Health Information)	Health information that can be linked to an individual patient. HIPAA requires specific controls on the collection, storage, processing, and transmission of PHI. All AI features in this strategy that involve patient data handle PHI.
SaMD (Software as a Medical Device)	A regulatory classification applied by the FDA to software that is intended to diagnose, treat, or prevent a disease or condition. Human-in-the-loop AI tools that support clinical judgment without making autonomous clinical decisions are generally outside current SaMD enforcement scope.
Vector Store	A database optimized for storing and retrieving high-dimensional embeddings (numerical representations of text or other data). Used in the AI Care Companion architecture to retrieve relevant clinical protocol passages in response to patient queries.

E. Interview List (Anonymized)

Role	Interview Format	Primary Topics Covered
Chief Executive Officer	90-minute structured + unstructured	Strategic context, board mandate, competitive landscape, AI thesis alignment, capital appetite
Chief Financial Officer	60-minute structured	Unit economics by program line, EBITDA bridge, capital allocation history, Series C terms overview
Chief Medical Officer	90-minute structured	Clinical protocol status, AI augmentation appetite, clinical safety requirements, governance preferences
Chief Operating Officer	60-minute structured	Partner reporting workflow, operational bottlenecks, change management capacity, process documentation status
Head of Engineering	90-minute structured + technical review	Platform architecture, integration depth, team capacity, AI engineering experience, build timeline estimates
Head of Data Science	90-minute structured + data review	Dataset assessment, prior AI experiment review, model development capacity, Databricks familiarity
Head of Clinical Operations	60-minute structured	Partner reporting process detail, NP/RD workflow documentation, operational AI readiness
VP Business Development	60-minute structured	Partner acquisition pipeline, partner reporting friction, new partner objection analysis
Board Member A (Healthcare Operating Advisor)	60-minute structured	AI strategy expectations, board decision criteria, competitive threat assessment, investment appetite

Role	Interview Format	Primary Topics Covered
Board Member B (Lead Investor Representative)	60-minute structured	Investment thesis alignment, Series C context, AI strategy board presentation requirements
Nurse Practitioners (4)	45-minute semi-structured each	Workflow documentation burden, between-appointment patient engagement, AI tool receptivity, session documentation time
Registered Dietitians (3)	45-minute semi-structured each	Metabolic and Cardiac program workflow, patient engagement patterns, documentation practices, AI augmentation concerns