KYTHERA
DECIPHERING HEALTHCARE

WHITE PAPER

# The Intelligent Strategy System

## How AI Agents and Open Claims Data Are Redefining Healthcare Strategy

***A Strategic Guide for Healthcare Leadership***


**A. Robert Grant**
*Chief Commercial Officer*
*Kythera Labs*

**Glynn Dennis, PhD**
*Chief Science Officer*
*Kythera Labs*


## Executive Summary


Health systems operate in an environment of compressed margins, intensifying competition, and exponential data growth. Yet the strategic planning functions responsible for navigating this environment remain constrained by manual workflows, fragmented data sources, and analytics cycles measured in weeks rather than hours. A new class of technology - AI agents equipped with composable analytical skills, connected to normalized open claims data on enterprise-grade platforms - is poised to fundamentally change how healthcare business leaders make decisions.

This white paper examines how the convergence of agentic AI, open claims data, and modern data intelligence platforms creates a step-change in strategic planning capability. It introduces the Kythera approach: purpose-built AI agent skills layered on normalized, interoperable healthcare data lakes powered by Databricks, enabling strategy teams to compress multi-week analyses into single-session interactions. For business development and strategic planning leaders evaluating where to invest, this paper provides the evidence, framework, and business case for operationalizing AI-driven strategy at the market level.





KYTHERA

| 40% | $27.7B | 2-3wks |
|---|---|---|
| of hospitals operating at a loss (2024) | healthcare claims management market | typical strategic analysis cycle |

*Sources: Strata Decision Technology, 2024; Mordor Intelligence, 2025; Kythera Labs internal benchmarking*

## The Strategic Intelligence Gap in Healthcare

### A Market Under Pressure

The U.S. healthcare industry faces a convergence of financial, operational, and competitive pressures that demand faster, more precise strategic decision-making. According to Kaufman Hall, hospital operating margins stabilized in 2024 at a median of approximately 4.9%, yet this headline figure masks deep disparity. Strata Decision Technology reported that 40% of hospitals were operating in the red in mid-2024. Moody's reported that median operating margins for not-for-profit hospitals improved from 0.5% in 2023 to 1.5% in 2024, better but still well below pre-pandemic norms. More than 700 rural hospitals face closure risk, according to the Center for Healthcare Quality and Payment Reform.

At the same time, expenses continue to climb. In 2025, hospital drug costs rose 9.3% year over year, total non-labor expenses grew 5.0%, and expenses per physician FTE approached $1.1 million annually. With expenses rising at roughly 6% annually and revenue increasing at only 3%, margin compression is structural rather than cyclical.

The longer-term trend reinforces this reality. Median operating margins declined from stable pre-pandemic levels of 2.1% to 2.6% in 2019 to a historic low of negative 0.5% in 2022, recovering to approximately 1.3% by the end of 2025. While improvement is evident, margins remain roughly one percentage point below the 3% level often cited as necessary for sustainable reinvestment and long-term financial stability.

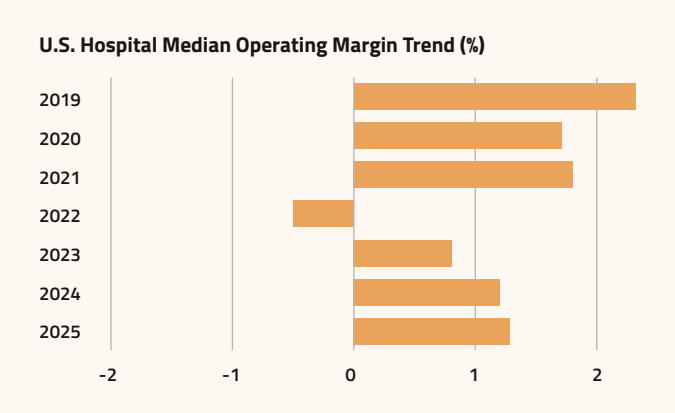

### The Data Paradox

Health systems generate enormous volumes of data across claims, encounters, referrals, financial transactions, and operational metrics. Yet strategy teams remain data-rich and insight-poor. According to Health Catalyst research, report development consumes the majority of healthcare analysts' time, leaving little bandwidth for actual analysis and strategic consulting. As report volumes grow, unmet demand for analytical capacity compounds.

**80%**

**of time spent on data preparation**

**20%**

**on actual insight generation**

The root causes are well documented: siloed EHR systems, inconsistent data definitions across facilities, fragmented claims data from multiple payers, and manual ETL processes that introduce latency and error. A typical strategic analysis - such as evaluating physician referral leakage across a multi-county service area - might take a team two to three weeks of pulling data from disparate systems, cleaning it, reconciling definitions, and building presentations. By the time the analysis is complete, the underlying conditions may have already shifted.

> The organizations that move first to operationalize these capabilities at the market level will have a structural advantage in speed, precision, and decision quality."
>
> — *Rob Grant, Kythera Labs*

## Current Solutions and Their Limitations

### Legacy Business Intelligence

Most health systems rely on a combination of enterprise BI tools (Tableau, Power BI, Qlik), data warehouses with scheduled ETL jobs, and consulting engagements for strategic analysis. These approaches served the industry well in an era of slower competitive cycles, but they are fundamentally mismatched to the velocity of modern healthcare strategy.

### Static Dashboards

Traditional dashboards answer predetermined questions with predetermined data structures. They cannot follow a chain of reasoning across datasets, adapt to novel questions, or synthesize findings into actionable recommendations. When a business leader asks a question that falls outside the dashboard's scope, the request enters a development queue that may take weeks to fulfill.

### Manual Claims Analysis

Open claims datasets from vendors like Definitive Healthcare, IQVIA, and state all-payer claims databases (APCDs) provide valuable market intelligence. However, working with these datasets typically requires specialized analysts who can navigate complex schemas, reconcile provider identities across sources, and apply business logic specific to the strategic question at hand. The process is linear, labor-intensive, and difficult to scale.

### Consulting Engagements

External consultants from firms like Sg2, Chartis, and Advisory Board bring deep domain expertise but operate on engagement timelines measured in weeks or months, at price points that limit their use to high-stakes decisions. The strategic planning function needs this caliber of insight on a daily basis, not quarterly.

**Table 1: Strategic Analysis Capability Comparison**

| Capability | BI Dashboards | Manual Claims | Consulting | AI Agents |
|---|---|---|---|---|
| **Time to Insight** | Hours–Days | 2–3 Weeks | 4–12 Weeks | **Minutes** |
| **Adaptability** | Low | Medium | High | **Very High** |
| **Cross-Domain Analysis** | Limited | Manual | Strong | **Automated** |
| **Cost per Analysis** | Low | Medium | Very High | **Low** |
| **Scalability** | Medium | Low | Low | **Very High** |

## The Accelerating Decision Clock: More Scenarios, Less Time, Less Money

The limitations of legacy solutions are not merely inconvenient - they are becoming existentially dangerous as the pace of healthcare strategic decision-making accelerates beyond the capacity of traditional tools and human-only workflows.

### The Volume of Decisions Is Expanding

A decade ago, a health system's corporate development team might evaluate two or three acquisition targets per year and conduct one major service line expansion study. Today, the same team faces a fundamentally different workload. Physician practice consolidation is accelerating - with 78% of physicians now employed by hospitals, health systems, or corporate entities - meaning the pool of available independent targets is shrinking and the competitive window for no acquiring the remaining practices is compressing. Simultaneously, the shift to outpatient care is multiplying the number of site selection decisions, payer contract renegotiations, and service line repositioning analyses that must occur in parallel. A single mid-size system might now run 10–15 distinct strategic evaluations in a year. A large IDN might run 30 or more.

Each of these decisions requires its own data assembly, market definition, competitive analysis, financial modeling, and strategic synthesis. The number of scenarios that must be modeled for each decision is also growing: What if we acquire Practice A but not Practice B? What if a competitor enters this market? What if payer reimbursement shifts by 5%? What if population growth in this zip code underperforms projections? Legacy tools force strategy teams to answer these questions one at a time, sequentially, with manual effort at every step.

### The Timeline for Each Decision Is Shrinking

Competitive dynamics are compressing decision timelines from quarters to weeks. When a private equity-backed platform announces its entry into a market, the health system's window to preemptively acquire key feeder practices may be measured in days, not months. When CMS releases new reimbursement rates or a competitor announces a new ambulatory surgery center, the strategic implications must be modeled immediately - not after a six-week consulting engagement. Deal value in healthcare M&A surged 56% in the first half of 2025 (KPMG), and 61% of investors expect further acceleration. The organizations that can evaluate, score, and act on opportunities fastest will capture the most strategic value.

Yet the tools available to strategy teams have not kept pace. A 4–12 week consulting engagement is not a viable response to a market that moves in 4–12 days. A dashboard that takes two weeks to reconfigure for a new question is not a viable tool for a business leader who needs an answer by Friday. The time pressure is structural, not situational - it reflects permanent changes in competitive velocity, regulatory cadence, and capital deployment speed.

**Budgets Are Tightening While Demand Is Rising**

The financial pressure on health systems is intensifying: 40% of U.S. hospitals are operating at a loss (Kaufman Hall, 2024), and more than 30% of rural hospitals are at risk of closure. Strategy and corporate development budgets are not growing proportionally to the volume of decisions these teams must support. Organizations cannot afford to commission a $200K–$500K consulting engagement for every acquisition evaluation or site selection study. They need the same depth of analysis at a fraction of the cost and time - and they need it repeatedly, not once.

The math is unforgiving: if a system needs to run 15 strategic evaluations per year and each consulting engagement costs $150K–$300K, the annual cost of externally sourced strategy intelligence alone approaches $2.25M–$4.5M - a figure that few mid-size systems can sustain and that even large IDNs are scrutinizing. The market is demanding more decisions, at higher quality, faster, and at lower cost per decision. This is not an incremental improvement

## The Decision Velocity Gap

Health system strategy teams face a widening gap between the volume and velocity of decisions required and the capacity of their current tools to support those decisions. Ten years ago, a strategy team might evaluate 3 acquisitions and 1 site selection per year with a 3-month timeline for each. Today, the same team must evaluate 15+ strategic scenarios across acquisitions, site selections, service line expansions, and payer renegotiations - often with decision windows measured in weeks. Legacy solutions were built for the old cadence. The market now demands a new one.

## The Critical Gap: From Information Delivery to Business Action

There is a deeper structural problem with every solution in the current landscape - one that transcends speed, cost, and scalability. Every existing tool, from BI dashboards to consulting engagements, stops at the same point: they deliver information and leave humans to do the work of converting that information into a decision and a business action.

A dashboard shows de that market share in orthopedics is declining in a specific geography. It does not tell the strategy team what to do about it. A claims analysis reveals that 12% of referrals are leaking to a competitor. It does not score the three practice acquisitions that would recapture the highest-value portion of that leakage, rank them by strategic fit, estimate the revenue impact, and produce a board-ready recommendation. A consulting engagement delivers a 150-page strategic plan. It does not continuously monitor whether the assumptions underlying that plan remain valid or alert the executive team when conditions change.

The entire existing solution landscape operates on a model where technology and services produce inputs to human judgment - data, visualizations, reports, benchmarks - and the human must then:

1. **Synthesize** findings across multiple data sources and analyses into a coherent picture
2. **Evaluate** trade-offs between competing priorities, resource constraints, and strategic objectives
3. **Formulate** a specific recommendation with supporting rationale
4. **Validate** that recommendation against additional data, scenarios, and sensitivities
5. **Package** the recommendation into a format suitable for executive or board-level review
6. **Act** on the recommendation - initiating a letter of intent, approving capital expenditure, restructuring a service line, or renegotiating a payer contract

Each of these steps introduces delay, cognitive load, and the potential for error or bias. Collectively, they represent the information-to-action gap - the distance between having data and doing something with it. In an environment where decision volume is rising, timelines are shrinking, and budgets are tightening, this gap is the single largest source of strategic underperformance in health system corporate development.

**Table 2: Where Current Solutions Stop vs. Where Health Systems Need to Go**

| Stage | Current Solutions Deliver | What Health Systems Actually Need |
|---|---|---|
| **Data Assembly** | Raw data feeds requiring analyst assembly | Pre-integrated, cross-source data environments automatically reconciled and linked |
| **Analysis** | Dashboards and reports answering pre-defined questions | Dynamic, multi-scenario analysis that follows chains of reasoning across datasets and adapts to novel questions in real time |
| **Synthesis** | Visualizations and summaries that humans must interpret | Composite scores and ranked recommendations that synthesize multiple signals into clear priorities |
| **Recommendation** | Not provided - left to the business team | Specific, confidence-weighted recommendations with supporting rationale, trade-off analysis, and sensitivity testing |
| **Validation** | Manual scenario modeling (if time permits) | Automated scenario generation and stress-testing across dozens of variable combinations in minutes |
| **Action Enablement** | Not provided - left to the business team | Board-ready output packages, deal scorecards, LOI-supporting analyses, and continuous monitoring that triggers alerts when conditions change |

The next generation of strategic intelligence for health systems must move beyond delivering information. It must deliver decisions - or, at minimum, decision-ready outputs that collapse the information-to-action gap from weeks to hours. This means composite scoring that synthesizes multiple data signals into a single actionable grade. It means ranked recommendation lists that tell a strategy team not just what the data says, but what it implies they should do. It means automated scenario modeling that tests a recommendation against dozens of "what if" conditions before a human ever reviews it. And it means continuous monitoring that does not wait for the next quarterly review to surface a change in market conditions - it alerts the right person, with the right context, the moment a strategic assumption is invalidated.

The organizations that close this gap first - that move from information provision to action enablement - will make better decisions faster, at lower cost, and with greater confidence. Those that do not will find themselves perpetually one step behind competitors who can.

## The Kythera Approach: AI Agents on Normalized Healthcare Data

Kythera Labs has spent years building the foundational data infrastructure required to make AI-driven healthcare strategy a reality - not as a prototype, but as an enterprise-grade capability. The approach rests on three interconnected pillars: normalized data, composable agent skills, and an enterprise intelligence platform.

### Pillar 1: Normalized, Interoperable Data Lakes

**Why Normalization Is the Prerequisite for AI-Native Strategy**

AI agents are only as effective as the data they reason over. A large language model or scoring algorithm that ingests healthcare data directly from source systems - raw claims feeds, unreconciled provider directories, inconsistently coded encounters, fragmented payer classifications - will produce outputs that are unreliable at best and dangerously misleading at worst. The fundamental barrier to deploying AI in health system strategic planning is not the sophistication of the models. It is the readiness of the data.

Most healthcare data exists in a state that is functionally hostile to AI reasoning: provider identities fragment across sources, the same procedure carries different codes depending on the billing system, payer classifications vary by state and data vendor, and geographic hierarchies are inconsistently structured. An AI agent asked to score acquisition targets across a multi-state market cannot produce a reliable composite grade if the underlying data treats the same physician as three different entities, maps the same CPT code to different service line categories, or fails to reconcile commercial payer data with Medicare claims at the patient level. The data must be made AI-ready before any intelligent system can act on it.

### Kythera's Normalization Layer: Building the AI-Ready Data Foundation

Kythera's data normalization layer is not a generic ETL pipeline or a data warehousing exercise. It is purpose-built infrastructure that transforms raw, fragmented healthcare data into a unified, AI-ready schema - a data environment specifically structured so that AI agents, scoring models, and automated reasoning systems can traverse it reliably, at speed, and without the brittle manual joins and transformation pipelines that have historically made cross-domain analysis so painful.

The normalization layer ingests data from multiple sources - open claims databases, FHIR-compliant clinical feeds, CMS public use files, payer-reported encounters, demographic datasets, and facility-level operational data - and performs four critical transformations that elevate raw data into an AI-ready substrate:

**Table 3: The Four Transformations That Make Healthcare Data AI-Ready**

| Transformation | What It Does | Why It Matters for AI |
|---|---|---|
| **Provider Identity Resolution** | Reconciles provider records across claims, directories, NPPES, and CMS files to create a single, authoritative identity for each practitioner, practice, and facility - regardless of how many source systems reference them under different IDs, names, or NPIs | AI agents scoring acquisition targets or mapping referral networks must treat each provider as a single entity. Without identity resolution, scoring models double-count volumes, misattribute referrals, and produce unreliable composite grades. Resolution is what makes provider-level AI reasoning possible. |
| **Coding Taxonomy Standardization** | Maps all procedure, diagnosis, and encounter codes to a harmonized taxonomy across ICD-10, CPT, HCPCS, and DRG systems - resolving version conflicts, local code extensions, and payer-specific coding variations | AI models that analyze service line opportunity, competitive capacity, or payer reimbursement must compare like to like. Standardization gives AI a consistent analytical vocabulary. |
| **Payer Classification Harmonization** | Normalizes payer identifiers and plan types into a unified classification schema that reconciles commercial, Medicare, Medicaid, managed care, and self-pay categories across data vendors, state reporting systems, and facility billing platforms | Payer mix is a critical input to acquisition scoring, revenue modeling, and rate benchmarking. AI agents must distinguish between Medicare Advantage and traditional Medicare, between Medicaid managed care and fee-for-service - consistently, across every data source Harmonization makes payer-aware AI analysis reliable. |
| **Geographic Hierarchy Structuring** | Organizes all location data into a consistent hierarchy from national to state, county, zip code, and census tract levels - linked to ESRI demographic projections, economic indicators, and provider catchment areas | Every module in the Kythera platform depends on geographic precision. An AI agent modeling population growth in a candidate market must seamlessly traverse from zip-code-level demographic data to county-level facility capacity to state-level regulatory context. Structured hierarchies make geographic AI reasoning seamless. |

**The Result: A Data Lake That AI Can Reason Over**

The output of Kythera's normalization layer is not simply a cleaner version of the same data. It is a fundamentally different kind of data asset - one where any query can traverse claims, encounters, demographics, financials, and operational metrics without manual intervention, and where AI agents can chain multi-step reasoning across domains with confidence in the consistency of the underlying information.

This distinction is critical because it redefines the value proposition of Kythera's data. Raw healthcare data, no matter how comprehensive, is an input that requires substantial human effort to operationalize. An AI-ready data lake is an infrastructure layer that enables a new class of automated, intelligent applications - from composite acquisition scoring to dynamic site selection to real-time competitive monitoring - that were previously impossible because the data could not support the reasoning required.

## The Ai-Readiness Multiplier

Kythera's normalization layer does not just clean data - it elevates data from a raw commodity into AI-ready infrastructure. This transformation is what makes every downstream module possible: acquisition scoring requires resolved provider identities; demand forecasting requires structured geographic hierarchies; payer rate benchmarking requires harmonized payer classifications; service line analysis requires standardized coding taxonomies. Without this foundation, AI agents produce unreliable outputs. With it, they produce the composite scores, ranked recommendations, and confidence-weighted decisions that close the information-to-action gap.

The normalization layer is not a feature of the product. It is the product's foundation - and it is the single most defensible technical moat in the Kythera platform.

**Why Competitors Cannot Replicate This Easily**

Building an AI-ready healthcare data lake is not a software configuration exercise. It requires deep domain expertise in healthcare coding systems, payer structures, provider organizational hierarchies, and the idiosyncrasies of dozens of data sources - combined with years of iterative refinement as edge cases, data quality issues, and source-specific anomalies are identified and resolved. Kythera's normalization layer reflects this accumulated expertise across 9.7 billion claims and 310 million patients. Competitors who attempt to layer AI agents on top of raw or lightly processed data will face the same reliability failures that have historically plagued healthcare analytics: misattributed volumes, fragmented provider identities, inconsistent geographic mapping, and payer mix calculations that do not reconcile across sources.

The normalization layer transforms Kythera from a data vendor into an AI-native decision intelligence platform. The data is not merely available - it is structured for machine reasoning, which means every AI agent, scoring model, and automated workflow built on top of it inherits the reliability and consistency of the foundation beneath it. This is the infrastructure that makes the transition from information delivery to business action possible.

## Pillar 2: Databricks Data Intelligence Platform

Databricks provides the compute and governance layer that makes this architecture enterprise-ready. Unity Catalog ensures data lineage and access controls satisfy HIPAA requirements. Delta Lake provides ACID-compliant storage that supports both batch and streaming workloads. The lakehouse architecture means that the same data can serve both the AI agent's real-time reasoning and the organization's traditional BI and reporting needs without duplication or synchronization overhead.

### Dialog with Data

The critical difference is that these aren't static dashboards or canned reports. The agent reasons over the data in real time, follows the logic chain wherever it leads, and produces outputs tailored to the specific strategic question being asked.

## Pillar 3: Composable AI Agent Skills

### How LLMs and Skills Work Together

A large language model is a reasoning engine. It can interpret natural language, synthesize information, evaluate trade-offs, and generate structured outputs. But on its own, an LLM has no access to healthcare data, no ability to query claims databases, and no mechanism to execute the multi-step analytical workflows that corporate development decisions require. It can think, but it cannot do.

Skills are what give an LLM the ability to act. The open-standard protocol for Agent Skills - a simple, open format for giving agents new capabilities and expertise (agentskills.io) - provides the analytical framework that connects an LLM's reasoning capabilities to real data sources, domain-specific business logic, and structured outputs - in order to take real business actions on the user's behalf. Originally developed by Anthropic and released as an open standard, the Agent Skills format has been adopted by a growing ecosystem of AI development tools and agent products. A skill is a purpose-built capability - a folder of instructions, scripts, and resources - that performs a specialized, multi-step task autonomously: it can query a FHIR-compliant data source, interpret the clinical and administrative records it returns, apply healthcare-specific business logic, and produce a structured, decision-ready output - all without a human manually stitching those steps together.

The relationship is symbiotic. The LLM provides the reasoning layer - it can perceive user intent, plan a custom workflow, and determine which analytical steps are needed to answer a question or complete a task. Skills provide the execution layer - the concrete analytical capabilities that the LLM selects, sequences, and chains together to accomplish that task. When a strategic planning leader asks "Which independent cardiology practices in our market should we acquire first?", the LLM interprets the intent, decomposes it into analytical steps, and then selects and chains the specific skills needed to answer it: referral network analysis, payer mix evaluation, competitive risk scoring, physician alignment assessment, and revenue impact projection - orchestrated as a single, coherent workflow.

## The Agent Architecture

**Skills (Execution Layer):** Purpose-built analytical capabilities - packages of instructions, scripts, and domain knowledge - that query data, apply business logic, and produce structured outputs

**Agent (LLM + Skills):** Selects and chains multiple agents, skills, and workflows to take real business actions on the user's behalf - from scoring acquisition targets to generating board-ready recommendations

*The open-standard Agent Skills protocol makes this architecture composable, extensible, and interoperable - new skills can be added to the library without re-engineering the agent, the LLM dynamically discovers and chains them based on the task at hand, and the same skills work across any skills-compatible agent product.*

### The Agent Skills Open Standard: Why It Matters

The Agent Skills format (agentskills.io) addresses a critical challenge in enterprise AI deployment: how to give agents reliable, specialized capabilities without locking into a single vendor's proprietary framework. The specification defines a portable, version-controlled format - skills are simply folders containing a SKILL.md instruction file, supporting scripts, and domain resources - that any compatible agent can discover and use.

For Kythera's healthcare customers, the open standard delivers three strategic advantages:

- **Portability and interoperability.** Kythera's healthcare-specific skills - referral network analysis, payer mix optimization, service line demand modeling, and others - are built once and deploy across multiple agent products. Health systems are not locked into a single AI vendor; they adopt capabilities that work across their evolving technology stack.
- **Enterprise knowledge capture.** Organizations can author their own skills to encode institutional knowledge - market-specific business rules, proprietary scoring criteria, internal workflow requirements - in portable, version-controlled packages that persist across personnel changes and can be audited for compliance.
- **Ecosystem extensibility.** As the Agent Skills ecosystem grows, new capabilities from third-party skill authors - regulatory compliance checks, clinical quality benchmarking, real estate market analysis - can be added to a health system's agent without custom engineering, expanding the platform's value over time.

### Why Composability Changes the Game

What makes skills transformative is not any individual capability in isolation. It is composability - the ability of the agent to chain multiple skills together, in sequences that are dynamically determined by the question being asked, to accomplish tasks that previously required an analyst, a data engineer, and a domain expert working in sequence over days or weeks.

The agent reasons through each step, decides which skill to invoke next, evaluates the output, adapts its approach based on what it finds, and chains the results forward into the next analytical step - operating at machine speed across enormous datasets. A workflow that would take a human team two weeks of manual analysis is completed in minutes, with full transparency into the reasoning chain and the data supporting each conclusion.

This composability is what bridges the gap from information delivery to business action. A non-composable system can answer a single, pre-defined question (e.g., "What is our market share in orthopedics?"). A composable agent can pursue a multi-step strategic objective (e.g., "Identify the three practice acquisitions that would most improve our cardiovascular referral network, score them against our strategic priorities, estimate the revenue impact of each, and produce a board-ready recommendation with sensitivity analysis"). The first is information. The second is action enablement.

### Kythera's Healthcare-Specific Skill Library

Kythera has developed a library of healthcare-specific skills that encode deep domain knowledge - the business logic, data relationships, and analytical frameworks that strategy consultants carry in their heads, now embedded in executable, composable capabilities that the agent can invoke on demand.

**Table 4: Healthcare-Specific Agent Skills and the Business Outcomes They Enable**

| Skill | What It Does | Customer Outcome |
|---|---|---|
| **Referral Network Analysis** | Maps patient pathways across the full care continuum, quantifying referral volumes between providers, identifying leakage points, and tracing downstream revenue flows from any practice or facility in the market | A system losing $12M annually in referral leakage can identify the three specific practice acquisitions that would recapture 60–70% of that lost volume - with revenue projections attached to each target - in an afternoon instead of a quarter |
| **Payer Mix Optimization** | Evaluates payer composition against reimbursement benchmarks, Medicare/ commercial rate differentials, and Kythera's Remit Model data to quantify revenue quality | Before committing $15M to acquire a multi-site primary care practice, a team can confirm that the target's payer mix will actually improve system-wide revenue per encounter - or discover an unfavorable Medicaid concentration makes the deal a net negative |
| **Service Line Demand Modeling** | Projects 5- and 10-year service demand by procedure category, incorporating population aging, chronic disease prevalence, migration patterns, and site-of-care shift assumptions | A system evaluating a $40M ASC investment can see that orthopedic volumes are projected to grow 22% over the next decade while competitor capacity is at 85% - or that another market is the better investment |
| **Competitive Landscape Assessment** | Calculates market share by service line, tracks competitor facility openings and closings, measures specialist density per 100K population, and identifies markets where positioning is strengthening or eroding | A CEO can show that a competitor's new outpatient center has shifted 8 market share points in cardiology - and present a scored response plan with three acquisition and two de novo options ranked by strategic impact |
| **Physician Alignment Scoring** | Scores provider loyalty by quantifying referral concentration, procedure volume trends, payer panel overlap, and practice ownership stability | A CMO can see that 4 of their top 12 referral sources have declining alignment scores and are within a PE group's acquisition footprint - with preemptive retention or acquisition recommendations attached |
| **Geographic Market Sizing** | Defines and sizes any custom market geography by population, insured lives, condition prevalence, utilization rates, and provider capacity | A strategy team can compare the total addressable patient population for three different market definitions in real time, adjusting boundaries until they find the configuration that aligns with their service line priorities |
| **Revenue Impact Projection** | Models financial impact of a proposed strategic action by combining claims-derived volume data with payer-specific reimbursement rates, capacity assumptions, and market growth projections | A CFO evaluating a $25M physician group acquisition can see a 3-year revenue model showing $8.2M in incremental annual revenue - with downside and upside scenarios from $5.8M to $10.6M depending on referral capture and payer mix assumptions |

**What This Looks Like in Practice: Composed Scenarios**

The real power of composable skills becomes visible when the agent chains them into end-to-end workflows that directly map to the strategic decisions health systems make every day. Below are three scenarios that illustrate how skills compose into business actions - not reports, not dashboards, but the specific outputs that business development leaders, CEOs, or CFO need to move from question to decision.

### SCENARIO 1 "WE'RE LOSING CARDIOLOGY REFERRALS. WHAT SHOULD WE DO ABOUT IT?"

**User intent:** Diagnose the source of referral leakage in cardiology and recommend a corrective strategic action.

**Skills chained:** Referral Network Analysis → Competitive Landscape Assessment → Physician Alignment Scoring → Payer Mix Optimization → Revenue Impact Projection

**What the agent does:** Maps all cardiology referral pathways in the defined market, identifies which practices are sending patients to competitors and why (proximity, payer acceptance, physician relationships), scores the at-risk referral sources by strategic value and acquisition feasibility, evaluates the payer mix quality of each potential target, and models the revenue impact of acquiring the top-ranked targets versus alternative interventions (employment contracts, co-management agreements, de novo site placement).

**Business action delivered:** A ranked list of three recommended actions - one acquisition ($4.2M projected annual revenue recapture), one employment agreement ($1.8M), and one de novo urgent care placement ($2.6M) - with board-ready scorecards, sensitivity analysis, and a proposed sequencing timeline.

### SCENARIO 2 "WHERE SHOULD WE BUILD OUR NEXT AMBULATORY SURGERY CENTER?"

**User intent:** Identify the optimal location for a new ASC based on unmet demand, competitive density, population growth, and financial viability.

**Skills chained:** Service Line Demand Modeling → Geographic Market Sizing → Competitive Landscape Assessment → Payer Mix Optimization → Revenue Impact Projection

**What the agent does:** Projects surgical procedure demand across candidate geographies for the next 10 years, sizes each market by insured population and condition prevalence, maps existing ASC and hospital outpatient capacity by service line, evaluates payer reimbursement rates at each candidate location against market benchmarks, and models the 5-year revenue trajectory for a new ASC at each site under conservative, base, and optimistic scenarios.

**Business action delivered:** Five candidate locations ranked by composite score (demand growth × competitive gap × reimbursement quality × population trajectory), with the top site showing projected year-3 revenue of $11.4M at 72% outpatient orthopedic utilization, a dynamic opportunity map, and a capital investment recommendation ready for CFO review.

## SCENARIO 3 "A PE-BACKED GROUP JUST ENTERED OUR MARKET. WHAT'S OUR EXPOSURE?"

**User intent:** Assess the competitive threat posed by a new market entrant and recommend a defensive strategy.

**Skills chained:** Competitive Landscape Assessment → Physician Alignment Scoring → Referral Network Analysis → Geographic Market Sizing → Revenue Impact Projection

**What the agent does:** Analyzes the PE group's documented acquisition pattern (specialties, practice sizes, geographic preferences), identifies which providers in the system's referral network match that acquisition profile, scores each at-risk provider by strategic value (referral volume, downstream revenue contribution, payer mix quality), models the revenue loss if those providers are acquired by the competitor, and generates a defensive action plan with preemptive acquisition, employment, or partnership recommendations prioritized by ROI.

**Business action delivered:** A threat assessment showing $18.3M in annual downstream revenue at risk across 7 high-value referring practices, with a prioritized defensive playbook: 2 immediate acquisition targets ($6.8M combined revenue protection), 3 employment agreement candidates ($4.2M), and 2 practices where a co-management arrangement is the most cost-effective response ($3.1M) - all scored, sequenced, and packaged for CEO and board review within 48 hours of the competitive threat being identified.

### From Information To Business Action

In each scenario, the agent does not produce a report and wait for a human to figure out what to do. It perceives the user's intent, plans a custom workflow, selects and chains the appropriate agents, skills, and workflows, and delivers a specific, scored, actionable recommendation that a decision-maker can act on. The human remains in control of the final decision - but the analytical work that used to consume weeks of team effort, hundreds of thousands of dollars in consulting fees, and untold cognitive load is completed in minutes.

This is the transition from providing information to taking business action. The skills library is the engine that makes it possible.

# Use Cases and Proof Points

The following scenarios illustrate how AI agents operating on normalized claims data transform the work of healthcare strategy teams. These represent capabilities the Kythera platform is designed to enable at the market level.

## Referral Leakage Recovery

A business development leader tasks the agent with analyzing physician referral leakage patterns across a three-county service area. The agent cross-references those patterns against payer mix and procedure volume trends, identifies the highest-value recovery opportunities by specialty, and generates a prioritized outreach plan - all in a single session. This analysis, which traditionally takes a team two to three weeks of pulling data from disparate systems, is compressed to under an hour.

**Estimated Impact:** A mid-size health system recovering even 5% of leaked referrals in high-margin specialties like orthopedics or cardiology could recapture $3–$8 million in annual net revenue.

## Physician Recruitment Intelligence

A business development director preparing for a physician recruitment meeting has the agent pull the current competitive landscape for orthopedics in a target market, model the revenue impact of adding a subspecialist based on local demand signals and referral flows, flag potential cannibalization risks with existing medical staff, and draft a pro forma grounded in real claims and encounter data.

**Estimated Impact:** Data-driven recruitment decisions reduce failed placement risk, where each failed physician recruitment costs an estimated $500,000–$1 million in search fees, onboarding, and lost productivity.

## Market Entry Assessment

A service line leader evaluating geographic expansion asks the agent to assess population health trends, unmet demand by diagnosis category, competitive provider density, payer reimbursement dynamics, and facility capacity constraints - synthesized into a market entry recommendation with supporting data, segmented by zip code or census tract.

**Estimated Impact:** Avoiding a single poorly informed market entry decision - where capital expenditure for a new ambulatory facility can exceed $20–$50 million - represents substantial risk mitigation.

# Implementation Framework

Operationalizing AI-driven strategy requires a phased approach that builds confidence, demonstrates value, and scales systematically.

**Table 5: Kythera Implementation Framework**

| Phase | Activities | Deliverables | Timeline |
|---|---|---|---|
| **1. Discovery** | Data source inventory; strategic use case prioritization; stakeholder alignment | Assessment report; data integration roadmap; priority use case catalog | 4–6 weeks |
| **2. Foundation** | Data normalization and ingestion; Databricks lakehouse deployment; HIPAA compliance validation | Operational data lake; validated data quality metrics; security certification | 8–12 weeks |
| **3. Agent Deployment** | Skill configuration for priority use cases; agent testing with strategy team; output validation | Production-ready agent with 3–5 core skills; user training completion | 4–6 weeks |
| **4. Scale & Optimize** | Additional skill development; cross-departmental expansion; performance benchmarking | Full skill library; ROI documentation; continuous improvement playbook | Ongoing |

## CRITICAL SUCCESS FACTORS

**Executive Sponsorship:** Alignment between business leadership and IT (e.g., CIO) on data strategy and AI governance is essential for organizational adoption.

**Data Quality Investment:** Agent outputs are only as reliable as the underlying data. The normalization layer is not optional - it is foundational.

**Domain-Specific Skills:** Generic AI tools lack the healthcare business logic required for strategic planning. Skills must encode payer dynamics, regulatory context, and clinical nuance.

**Human-in-the-Loop Validation:** The agent is a force multiplier, not a replacement. Strategy professionals validate outputs, provide contextual judgment, and make final decisions.

# ROI Analysis

The return on investment for AI-driven strategic planning operates across three value dimensions: time recovery, revenue capture, and risk avoidance.

## Time Recovery

Health Catalyst research demonstrates that report development consumes the majority of healthcare analysts' time. By automating data aggregation, normalization, and initial analysis, AI agents can recover 60–80% of the time currently spent on data preparation. For a strategy team of five analysts averaging $120,000 in fully loaded compensation, this represents $360,000–$480,000 in recovered analytical capacity annually - capacity that can be redirected toward higher-value strategic consulting.

## Revenue Capture

Faster, more granular market intelligence enables earlier identification of growth opportunities. Referral leakage recovery alone can generate $3–$8 million annually for a mid-size system. Service line optimization informed by real-time demand modeling can improve procedure volumes by 5–15% in targeted specialties. Physician recruitment supported by claims-based demand analysis reduces failed placement costs.

## Risk Avoidance

Strategic decisions in healthcare carry enormous capital implications. Poorly informed market entry decisions, service line investments, and M&A evaluations can result in tens of millions in sunk costs. AI-driven analysis provides a more rigorous, data-complete foundation for these decisions, materially reducing the probability and cost of strategic error.

**Table 6: Estimated ROI for Mid-Size Health System (200–500 beds)**

| Value Dimension | Conservative Estimate | Aggressive Estimate |
|---|---|---|
| **Time Recovery** | $360K/year | $480K/year |
| **Revenue Capture** | $3M/year | $8M+/year |
| **Risk Avoidance** | $5M (per decision) | $50M+ (per decision) |
| **Total Year-1 ROI** | **5–10x** | **15–25x** |

# Conclusion

Healthcare strategy has reached an inflection point. The combination of financial margin pressure - with 40% of hospitals operating at a loss and expenses outpacing revenue growth - competitive disruption from non-traditional entrants, and the sheer complexity of modern healthcare markets demands a fundamentally different approach to strategic intelligence.

The technology to deliver that approach now exists. AI agents with composable, healthcare-specific skills - built on the open Agent Skills standard and operating on normalized open claims data within enterprise-grade platforms - can compress weeks of analysis into hours, surface insights that manual processes would miss, and provide the strategic planning function with an analytical capability that scales with the complexity of the decisions it supports.

The question for healthcare business leaders is no longer whether AI will transform healthcare strategy - it is whether your organization will be among those that capture the first-mover advantage or those that spend the next several years trying to catch up.

**Sources and References**


Advisory Board / Kaufman Hall. National Hospital Flash Report: 2024 Year-End Data. February 2025.
Agent Skills. A Simple, Open Format for Giving Agents New Capabilities and Expertise. 2025. https://agentskills.io
Chartis Group. Hospital Margins Trend Higher, But Looming Disruption Demands Strategic Decisions. 2025.
Health Catalyst. How to Build a Healthcare Analytics Team and Solve Strategic Problems. 2025.
Mordor Intelligence. Global Healthcare Claim Management Market. 2025. Valued at USD $27.72 billion.
Moody's Ratings. Annual Not-for-Profit and Public Healthcare Median Financial Report. 2024.
Strata Decision Technology. U.S. Hospitals See Financial Performance Stabilize Throughout 2024. January 2025.
Strata Decision Technology. Monthly Healthcare Financial Benchmarks. November 2025.
Vizient / Sg2. Trends to Drive Your Healthcare Strategy in 2025. 2024.
Center for Healthcare Quality and Payment Reform. Rural Hospitals at Risk of Closing. 2025.
Fitch Ratings. 2025 Outlook: Rising Expenses Will Pressure Hospital Finances. December 2024.


**Connect with Kythera.** Kythera is a data technology company that brings unprecedented clarity and structure to complex real-world healthcare data. Kythera's Wayfinder Technology Platform, supported by pre-configured pipelines, processing libraries, analysis tools and remastered datasets, helps Healthcare and Life Sciences organizations work with greater speed, scale and confidence. Learn more at www.kytheralabs.com.




KYTHERA™
DECIPHERING HEALTHCARE

346 Main Street
Franklin, TN 37064
www.kytheralabs.com