# Securing Patient Data in Healthcare: From Legacy APIs to Agent2Agent (A2A) Protocols

## A White Paper on Protecting Health Information and Enabling Compliant AI Collaboration for Providers

**Published by Backflipt (Xenovus, Inc), November 2025**

---

## Executive Summary

In the AI-enhanced healthcare ecosystem, providers face a critical challenge: providing personalized patient insights and predictive analytics while safeguarding protected health information (PHI) from irreversible breaches. Traditional API and Model Context Protocol (MCP) frameworks expedite data access but increase the risks of PHI exposure through public large language models (LLMs), HIPAA violations, and missed opportunities for secure advisory services.

This white paper examines a real-world case study in which a healthcare provider predicts patient readmission risks for chronic conditions in Latin America, utilizing a medical analytics firm's historical datasets. We compare the "before" state—legacy APIs resulting in elevated costs, breaches averaging $7.42 million, and operational silos—with the "after" state facilitated by Agent2Agent (A2A) protocols and Lumen from Backflipt.

Embracing A2A enables providers to manage secure, autonomous processes that provide actionable insights without requiring raw data transmission. This advancement strengthens data privacy, upholds HIPAA compliance, and generates revenue via customized predictive services. Healthcare entities achieve faster, traceable decisions, lessening in-house modeling loads and bolstering patient trust.

**Key takeaways:**

- Legacy systems spread PHI leaks at clinical speed, threatening compliance and trust.
- A2A protocols promote interoperable AI agent interactions, emphasizing outcomes over raw datasets for privacy-focused collaboration in diagnosis and scheduling.
- **Lumen** optimizes A2A deployment, enhancing security and scalability in regulated healthcare settings.

---

# Introduction

The rise of artificial intelligence (AI) is revolutionizing healthcare, from predictive diagnostics to personalized treatment and administrative automation. Providers rely on medical analytics firms for enhanced datasets on patient outcomes, disease patterns, and compliance metrics to inform care decisions. However, as AI agents integrate into electronic health records (EHRs) and workflows, standard data channels like APIs and MCPs pose severe risks to PHI privacy and intellectual property (IP).

Introduced by Google in April 2025 and established as an open standard through the Linux Foundation's Agent2Agent Protocol Project, A2A marks a revolutionary shift. It provides AI agents with a secure framework for sharing context, status, actions, and results—without relying on underlying data, memory, or tools. This is essential for healthcare, where proprietary patient histories drive competitive advantages under rigorous regulations like HIPAA.

This white paper analyzes legacy vulnerabilities through a healthcare-specific scenario and demonstrates how **A2A, enhanced by Backflipt's Lumen**, mitigates these risks. Based on protocol standards and industry analyses, it maps a route for providers to retain control, comply with HIPAA, and leverage AI for advancement.

---

# The Challenge: PHI Risks in Healthcare Data Sharing

Medical analytics firms assemble high-value datasets from patient records, treatment outcomes, and regulatory data, licensing them to providers via APIs or MCPs for smooth integration. MCPs improve speed—yielding query responses in seconds—but unintentionally heighten dangers.

**Core challenges include:**

- **PHI Leakage via LLMs:** Providers processing raw feeds in public LLMs for analytics inadvertently incorporate sensitive details into third-party weights, resulting in unrecoverable breaches; examples include clinicians using ChatGPT with patient data, which can lead to HIPAA fines.
- **Regulatory Breaches:** Unmanaged pipelines can trigger HIPAA violations, with AI-related exposures, such as misconfigured systems or unauthorized sharing, resulting in penalties and reputational harm; the average breach cost is $7.42 million.
- **Commoditization and Revenue Loss:** APIs and MCP reduce analytics to low-value conduits, overlooking premiums for tailored insights on emerging conditions.
- **Provider Inefficiencies:** Entities incur rising costs for internal modeling, hiring specialists, and developing custom predictions, often bypassing the expertise of analytics firms amid privacy concerns.

These problems escalate in innovative analytics for volatile populations, requiring secure extrapolation of historical data.

---

# Scenario: Predicting Patient Readmission Risks in Latin America

Imagine a leading hospital network expanding chronic care programs in Latin America to address rising diabetes and cardiovascular cases. It needs to forecast readmission risks—considering factors like treatment adherence, regional demographics, and comorbidity patterns. Medical analytics firms supply historical patient data (e.g., discharge summaries, outcome metrics) but lack direct insights on this emerging cohort due to its novelty.

The provider requires a risk model for resource allocation, care planning, and outcome improvement. In legacy setups, this involves pulling raw datasets for in-house adaptation—a pathway to inefficiencies and PHI risks.

---

## Before A2A: The Legacy API/MCP Model

In the standard approach, the provider accesses the analytics firm's data via APIs or MCP:

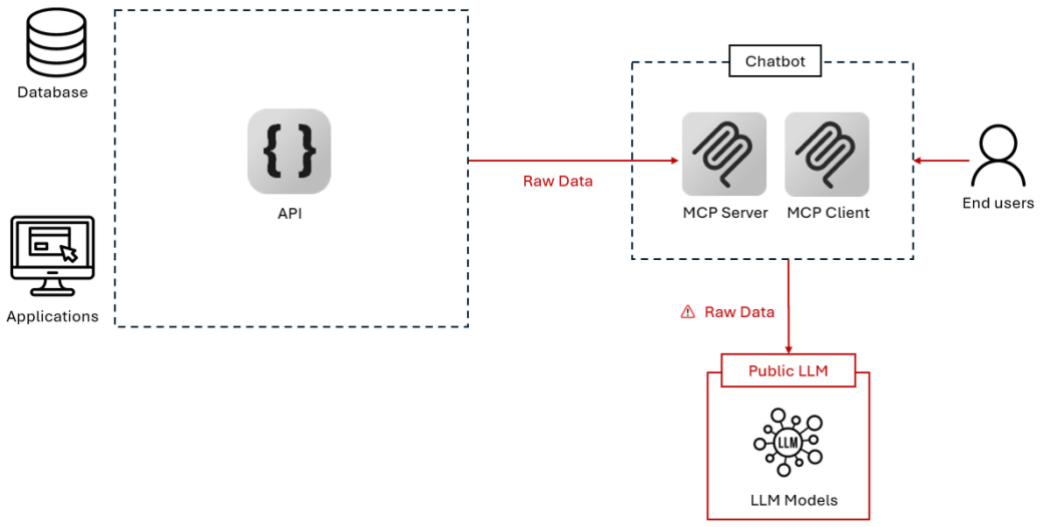


- **Process Flow:**
    1. Provider retrieves raw patient datasets (e.g., anonymized EHR extracts, outcome logs) via API/MCP.
    2. Inputs rows into public LLMs for initial risk scoring.
    3. Recruits external analysts to create custom readmission models.

- **Impact on the Provider:**

| Process Step | Impact |
|---|---|
| **Pull raw patient datasets via API/MCP** | Rapid access in seconds but sparks high modeling costs and exposure risks. |
| **Feed rows into public LLM** | Endangers PHI to breaches, IP theft, HIPAA violations, and fines over $7.42 million. |
| **Hire external analyst & build custom readmission model** | Accrues costs and risks non-standard models. |

- **MCP Advantage:** Accelerates data interaction—retrieval in seconds. **Greater Risks & Costs:** Mounting expenses, delays, security threats, and DIY inefficiencies outweigh the speed.
- **Impact on Medical Analytics Firms:**

| Exposure | Consequence |
|---|---|
| **Raw data enters public LLM** | IP trains 3rd-party models → permanent loss. |
| **No modeling involvement** | Miss upsell on readmission forecast service. |
| **API/MCP = commodity feed** | Revenue capped at data access fees. |

- **PHI Loss Math:** 1 PHI row → 1B-parameter model → No recall. Speed of MCP = Speed of PHI leakage.
- **Outcome (Before A2A):**

| Provider | Analytics Firm |
|---|---|
| **Fast data, slow insight; DIY modeling burden.** | PHI exposed globally; Lost control + lost revenue. |

**Legacy API/MCP = faster access, faster PHI death.** Speed without control leads to commoditized data and diminished future value.

---

# After A2A: Secure Collaboration with Lumen

A2A mitigates these issues by enabling AI agents to collaborate without requiring the sharing of raw data. As an open standard, A2A supports secure, interoperable communication where agents exchange goals, context, state, actions, and results. This "team-like" interaction ensures auditability and trust, preventing data leaks while supporting complex tasks like multi-agent diagnosis or billing.

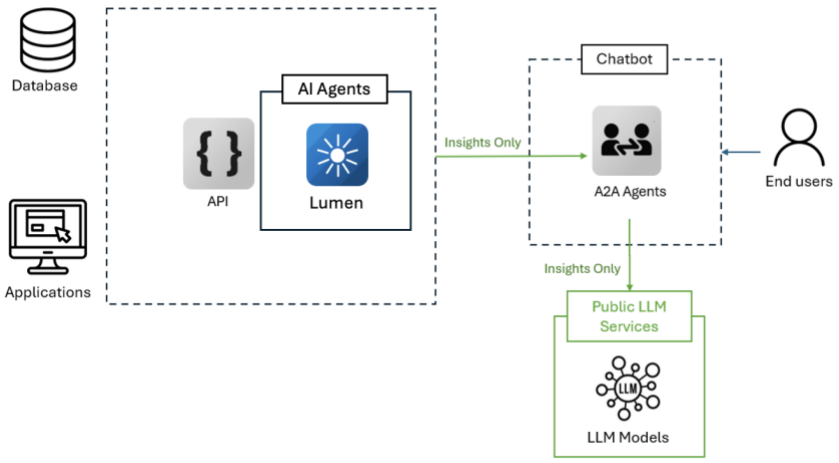


**Backflipt's Lumen** automates A2A agent creation for existing APIs, deploying seamlessly next to API gateways. Lumen integrates with secure managed file transfer (MFT) and business-to-business integration (B2Bi) environments, making it enterprise-ready for analytics firms in the healthcare sector.

## How A2A Works in This Scenario

1. **Deployment:** Lumen fronts the analytics firm's API, auto-generating an A2A agent per client or query type.
2. **Query Handling:** The provider's A2A agent sends a structured request (e.g., "Project readmission risks from chronic trends in Latin America") via the A2A protocol.
3. **Internal Processing:** The analytics firm's A2A agent computes the forecast using proprietary data and models, without exposing raw rows.
4. **Response Delivery:** Returns only insights (e.g., "LatAm readmission rate: 18%; Key factors: Adherence 65%"), maintaining full audit logs.

---

- **A2A Model Flow:**

| Process Step | Secure Flow |
|---|---|
| **Lumen fronts Analytics API** | Auto-spins A2A agent; shields raw patient data behind protocol. |
| **Provider A2A agent queries** | Sends structured request via A2A. |
| **Analytics A2A agent responds** | Computes insight internally; delivers only results. |

- **A2A Advantage:** Protocol ensures interoperability, auditability, and zero raw exposure—agents collaborate like teams, not data dumps. In healthcare, it enables modular systems for diagnosis and scheduling without privacy risks. **Lumen Edge:** Deploys in minutes; scales to multi-agent workflows.
- **Impact on the Provider:**

| Process Step | Impact |
|---|---|
| **Query via A2A agent** | Instant access to expert forecasts; no data pulls or security overhead. |
| **Receive structured insights** | Zero exposure to breaches, PHI theft, or compliance violations. |
| **Integrate ready projections** | Saves costs and time on modeling; taps industry expertise directly. |

- **A2A Speed:** Collaboration in seconds—far beyond MCP, with built-in trust. **Net Gain:** Accurate, compliant results without the DIY drag; supports privacy-preserving techniques like federated learning analogs.
- **Impact on Medical Analytics Firms:**

| Control | Consequence |
|---|---|
| **A2A gates all interactions** | Raw IP stays internal; trains only your models → permanent sovereignty. |
| **Protocol enables custom services** | Capture upsells on readmission modeling; bill for agentic insights. |
| **Beyond commodity APIs** | Revenue unlocks at premium tiers (e.g., $0.50+/query). |

- **Secure Value Math:** 1 A2A exchange → Auditable insight → Infinite reuse. Protocol speed = Locked PHI + Exponential revenue.
- **Outcome (After A2A):**

| Provider | Analytics Firm |
|---|---|
| **Seamless insight, full compliance; Expert accuracy, no burden.** | IP fortified, zero leakage; New service lines + deepened partnerships. |

**Legacy MCP = raw risk. A2A + Lumen = agentic gold.**

---

# Lumen vs. Alternatives: Why Lumen Wins

Lumen offers an **out-of-the-box A2A experience** that converts legacy APIs into compliant, monetizable agent endpoints in minutes—without coding, AI expertise, or disjointed workflows.

| Criteria | Backflipt Lumen | DIY/Open-Source (e.g., Langflow, LangGraph) | Native Cloud Tools (e.g., AWS Bedrock, Azure AI, Google Vertex) |
|---|---|---|---|
| **Core Focus** | **Instant API-to-AI transformation:** MCP servers & A2A agents in minutes, enterprise-hosted for control & monetization | Low-code AI workflows; manual prototyping | General AI agents; platform-specific integrations |
| **Speed to Market** | **Fastest:** Deploys in minutes via Swagger/gateways, shrink-wrapped instrumentation | **Slowest:** Manual build, test, deploy; needs expertise | **Moderate:** Fast prototyping, but custom dev for enterprise |
| **Enterprise Governance** | **Built-in:** CI/CD integration, automated consistency, unified dev portal | **Manual:** Governance from scratch; fragmented | **Component-Based:** Manual config for governance |
| **Ease for Non-Experts** | **High:** No-code POET for customization, no AI expertise needed | **Medium:** Visual UI, but coding for advanced features | **Variable:** Some no-code, but devs needed for complexity |
| **Scalability & Security** | **Enterprise-Grade:** Kubernetes auto-scaling, DevSecOps, compliance (HIPAA, GDPR) | **Custom:** Scaling/security depend on expertise | **Scalable:** Cloud-native, but custom app logic/security |
| **Total Cost of Ownership** | **Lowest:** 80–90% dev cycle reduction, minimizes talent needs, 2–3x revenue | **Highest:** High maintenance, dedicated teams | **High:** Low initial cost, but escalates with dev hours |

**Out-of-the-Box Advantage of Lumen:**

- **Zero-code agent generation** from existing APIs
- **Auto-instrumented A2A compliance & audit trails**
- **One-click deployment** next to API gateways
- **Built-in monetization layer** (per-insight billing, usage tiers)
- **No AI or agent framework expertise required** — non-technical teams launch in under 15 minutes

# Benefits of Transitioning to A2A

- **For Providers:** Maintain PHI control, lower breach risks, and monetize advanced services like AI-driven predictive care. A2A establishes providers as collaborative leaders rather than data silos.

- **For Medical Analytics Firms:** Utilize secure, opaque exchanges to prevent leaks while delegating tasks for deeper insights; interoperability reduces silos in diagnostics and compliance.
- **Broader Ecosystem:** Advances open protocols, driving AI innovation in healthcare—securely and scalable, addressing biases and privacy in data sharing.

Pioneer adopters report **30–50% revenue uplift** from premium insights and up to **80% reduction** in data exposure incidents.

---

## Conclusion

The age of unguarded data sharing is past. As AI agents expand in healthcare, providers must transition from legacy APIs to A2A protocols to protect their core assets. Backflipt's Lumen makes this shift effortless, enabling secure, intelligent collaborations that benefit all parties.

By adopting A2A, entities can "own the model and sell the future"—transforming risks into opportunities. We suggest piloting A2A integrations with Lumen for high-value scenarios like readmission forecasts to achieve immediate gains.