



The Case for Identity as Infrastructure in Health Tech:

The tech stack decision that
unlocks product innovation, AI,
interoperability, and growth



Introduction

Cracks in the foundation: The missing layer of a modern tech stack

Health tech companies are modernizing nearly every layer of the stack. Cloud infrastructure is table stakes. Data lakehouses, CRMs, digital front doors, and analytics platforms are standard operating investments. AI initiatives are funded, staffed, and on the roadmap. The pace of investment shows no signs of slowing.

But underneath all of it, quietly undermining operations and product development, is a gap many organizations haven't fully reckoned with: a solid identity foundation. We call that foundation identity intelligence.

Identity intelligence—the ability to consistently and accurately know who is who across systems, while continually managing and enriching that data—is still too often treated as a backend task handled through homegrown logic, a legacy Master Data Management (MDM) solution, or fragmented point solutions. In the early stages of a company's growth, a homegrown matching algorithm or basic open-source tool is a common choice. Duplicate records get cleaned up manually and edge cases get logged and revisited. But when the volume of data grows, what was once a manageable inefficiency becomes an enterprise constraint.

This lack of modernization is proving costly. When an organization cannot consistently identify a patient, a member, a user, every downstream workflow inherits that uncertainty. An AI model trained on fragmented identity data produces fragmented outputs. The interoperability initiative built on mismatched records creates noise. The analytics team draws conclusions from a population that includes dupes, ghosts, and gaps.

This white paper makes the case for **treating trusted identity as the foundation for every downstream investment including AI, interoperability, analytics, and digital experiences**. It's written for enterprise IT, data, product, AI, and analytics leaders in healthcare technology organizations who are ready to close the gap and build on a solid foundation and unlock identity intelligence. You'll learn about Opala, too—a company that's making all the right moves by treating identity as infrastructure.

According to Pew Charitable Trusts, patient matching accuracy can be as low as 80% within a single care setting and drops to just 50% among organizations that share the same EHR vendor.¹

The infrastructure shift:

Health tech has seen this pattern before

Health tech companies no longer build every foundational capability in-house. Cloud compute, data storage, security infrastructure, analytics engines—these are layers that health tech companies depend on completely and maintain directly almost not at all. The decision to outsource them wasn't made because they are unimportant. It was made because they are too important, and too complex, to recreate well while also building a differentiated product.

Health tech learned this lesson with cloud infrastructure the hard way. For years, organizations ran their own data centers, their own security teams, their own disaster recovery systems. The resistance to moving off that model was real. Then the economics and the evidence became impossible to ignore: cloud providers offered more security, more reliability, and lower total cost than any individual organization could match internally. The question shifted from “should we move to cloud?” to “why haven't we already?”

Identity is following the same arc.

Every health tech company needs to know who is who. The underlying problem is not unique to any single organization; it's universal. And yet many are still solving it independently, with homegrown logic, legacy platforms, or point solutions never designed to serve as enterprise infrastructure.

The most forward-thinking health tech companies are treating identity the way they treat cloud and data platforms — as a specialized layer that is too important to get wrong and too complex to build well while also building everything else.

The organizations that recognize this pattern early have a structural advantage over those still waiting to see how it plays out.



Why identity has become a strategic issue: The pressure is coming from every direction

Healthcare is now one of the most data-intensive industries on the planet. With a compound annual growth rate of 36%, healthcare data is expanding faster than manufacturing, financial services, and media combined.² Every new EHR integration, every wearable device, every claims feed, every digital health interaction adds to a data environment that is orders of magnitude more complex than it was a decade ago.

Five pressure points are driving health tech companies to revisit their identity management solutions.

1. Cloud and data platform modernization



Enterprise health tech companies are migrating to cloud-native architectures and modern data platforms, like Snowflake, Databricks, Azure, Google Cloud. These investments are made to unlock faster analytics, better data sharing, and scalable infrastructure. But when identity data migrated into these environments is fragmented, inconsistent, or poorly resolved, the modernization effort inherits the problem rather than solving it. A lakehouse built on uncertain identity data is not a foundation—it's a liability.

2. AI and analytics initiatives



According to McKinsey's Q4 2025 survey of healthcare leaders, 50% of healthcare organizations have now implemented generative AI, up from just 25% in late 2023.³ But AI does not fix fragmented identity. It amplifies it. AI is an unforgiving environment for poor identity data. Models trained on fragmented or duplicate records learn the wrong patterns. Populations used for risk stratification, care management, or predictive analytics that include identity errors produce outputs that cannot be trusted. And in healthcare, untrustworthy outputs are not just a product problem, they are a patient safety and liability problem. The organizations investing most aggressively in AI are discovering that identity infrastructure is not a prerequisite they can defer. It is the work that has to happen before the models can perform.

3. Acquisitions and platform expansion



Growth through acquisition is one of the most common strategies in health tech, and one of the most reliable triggers for an identity crisis. B2B healthcare IT deal volume rose 24% year-over-year in 2024, driven by private equity and buyers pursuing scalable interoperable technology platforms.⁴ And when two companies come together with two completely different approaches to identification, the value of that merger takes a hit. The cross-sell opportunity that justified the acquisition cannot be realized until the identity layer is reconciled. The unified platform the market was promised cannot be delivered until the organization knows, with confidence, that the records in both systems belong to who they say they do. Identity reconciliation is not a post-merger cleanup task. It is a growth enabler that sits on the critical path.



4. Interoperability requirements and ecosystem connectivity

The regulatory and commercial pressure to exchange data across systems, partners, and networks is intensifying. Payers and providers are expected to share data in near real time. Health tech companies are being evaluated on their ability to plug into broader ecosystems. Health information exchanges are only as valuable as the accuracy of the identity layer connecting them. When identity is fragmented, interoperability doesn't solve the problem—it scales it.



5. Customer experience

Patients and members are increasingly interacting with health tech platforms the way they interact with any digital consumer product. They expect the platform to know who they are, remember their history, and deliver experiences that feel consistent across every touchpoint. That expectation has only intensified as digital health adoption has grown. Telehealth, for example, remains significantly above pre-pandemic utilization levels,⁵ and every new digital touchpoint adds another moment where fragmented identity can break the experience.

When identity infrastructure cannot reliably connect a person's records across systems, those expectations go unmet. Re-enrollment friction, mismatched histories, and broken continuity of experience are not just annoyances. They erode trust, increase churn, and signal to the market that the platform is not ready to operate at scale.



The cost of fragmented identity:

What breaks when identity is not built for scale

For most health tech companies, the identity problem doesn't surface as a single catastrophic failure. It surfaces as persistent friction in different forms throughout the organization. Below are some of the common costs of not getting identity right.



Product and platform
velocity slows

Building on top of fragmented identity data is like building on top of shifting ground. The foundation looks stable until you try to move fast. Then the cracks show.

Engineering teams that should be shipping differentiated product capabilities instead spend cycles normalizing, reconciling, and stitching identities across systems. Data pipelines require constant maintenance to account for mismatches. New use cases that depend on a trusted, unified view of a person or provider take longer to stand up than the product roadmap anticipated. Integrating a new data source means revisiting identity logic that was never designed to be extensible.

The compounding effect is significant. Every sprint cycle spent on identity plumbing is a sprint cycle not spent on the capabilities that differentiate the product in the market. Over time, organizations that treat identity as something to build and maintain in-house are effectively taxing their own engineering capacity: paying an ongoing identity debt out of their product velocity budget.



AI and analytics become
less trustworthy

AI initiatives are only as reliable as the data they are built on. And identity is the layer that determines whether that data is complete, accurate, and connected enough to trust.

Duplicate records inflate population counts and skew model training. Fragmented identities create gaps in longitudinal data that algorithms interpret as clinical signals rather than data quality failures. Inconsistent identity resolution across source systems means that the same person may appear as two or three distinct individuals depending on which system is queried, and the model has no way of knowing.

The result is AI and analytics outputs that are difficult to validate and harder to act on with confidence. Risk stratification models built on fragmented populations misclassify patients. Predictive analytics tools surface recommendations that don't hold up under clinical scrutiny. Dashboards and reporting tools produce numbers that disagree with each other because they are drawing from identity layers that were never reconciled.

For health tech companies with AI on the roadmap—or already in production—this is not a future problem. Scaling AI without first resolving identity results in broad outputs that are unreliable.



Operations stay expensive and brittle

The operational cost of poor identity infrastructure is chronically underestimated—not because the costs are small, but because they are distributed across teams and budget lines in ways that make them hard to see as a single number.

Manual reconciliation consumes analyst and engineering hours that are accounted for as normal operating expense rather than identity debt. Exception handling queues grow as volume increases and the matching algorithm fails to keep pace. Identity-related support tickets—patients or providers mismatched in the system, records that don't resolve correctly, data that disagrees across platforms — become a recurring feature of the operations backlog rather than an anomaly to be fixed.

Legacy MDM platforms and homegrown solutions add a different dimension of cost. As these systems age, they require increasing IT investment to maintain, patch, and operate at scale. Institutional knowledge about how the system was configured concentrates in a small number of people. Consulting engagements to support aging infrastructure become a permanent line item. And when a forced migration arrives, as it has for many organizations with a legacy MDM solution, the true cost of staying with a subpar solution finally becomes visible all at once.

The total cost of ownership of a legacy or homegrown identity system is significantly larger than the software license suggests. It includes compute infrastructure, IT overhead, engineering maintenance, manual reconciliation labor, support staff, and consulting fees, most of which are invisible until someone adds them up.



Growth initiatives get harder

Identity infrastructure that was adequate for a single product line at a single scale becomes a constraint the moment the business tries to grow beyond it.

When two companies merge, they bring two different approaches to identification, two different matching algorithms, and two patient or customer populations that overlap in ways neither system can yet see. The cross-sell and expansion opportunity that justified the acquisition sits behind an identity reconciliation problem that can take months or years to resolve if the underlying infrastructure was not built for it.


Organic growth creates its own version of the same problem. As health tech platforms expand across new markets, new product lines, and new client environments, the identity layer has to scale with them. Connection-by-connection interoperability—where each new integration requires its own identity reconciliation effort—does not scale. It creates a ceiling on how fast the platform can grow and how many simultaneous integrations it can support without accumulating identity debt that slows every subsequent initiative.

The organizations that grow fastest are the ones that have resolved this constraint before it becomes a bottleneck.

What identity as infrastructure actually means: A trusted identity layer for systems of record, engagement, and insight

The ideal state is clear: a single, accurate, continuously resolved view of every person and organization across the enterprise — one that every system, every product, and every team can trust without having to recreate it themselves. That is what identity infrastructure means in practice. It is handled as a shared foundational layer across the enterprise, not recreated inside every product, workflow, or integration.

That shared layer sits between the systems that hold data and the systems that use it.



Without it, every team crossing that boundary is on their own. Data scientists build workarounds into their pipelines. Product engineers write reconciliation logic into their features. Operations teams hire people to manage the exceptions. Each solution is local.

When identity is handled as shared infrastructure rather than a per-team problem, the downstream effects are significant and compound over time. New data sources integrate faster because the identity layer is already configured to receive them. There is no per-source tuning, no algorithm reconfiguration, no months-long implementation cycle every time the platform expands.



AI and analytics become more reliable because models and reports are drawing from a population that has been accurately resolved. Duplicates don't inflate counts. Gaps don't create false signals. Data scientists can trust what they're working with before they build, not after they've already shipped.



Patient and customer experiences get cleaner because the platform knows who someone is across every touchpoint. Their history follows them, their preferences are recognized, and they don't have to re-identify themselves every time they interact with a new part of the product.



Technical debt decreases over time rather than accumulating. The identity reconciliation logic that currently lives in a dozen different places—in pipelines, in product code, in manual processes—gets consolidated into a single layer that is maintained once and trusted everywhere.



And when the organization acquires a new company or launches a new product line, **activation is faster because the identity infrastructure is already in place.** The new platform doesn't need to build its own identity logic from scratch. It plugs into what already exists.

This is what health tech companies that have made the infrastructure shift are actually gaining: not just cleaner data, but a compounding structural advantage that gets stronger every time they add a new system, a new product, or a new market to the platform.

What identity as infrastructure looks like in practice: **Meet Opala**

Opala didn't stumble into the identity problem. Its founding team came out of Premiera Blue Cross, where they had spent years building interoperability infrastructure from the inside. They understood firsthand how hard identity was to get right, and how much everything else depended on it. "It was clear to us that if you don't master identity first, you're just perpetuating erroneous data. You're spreading risk faster than you're gaining any benefit, and that's a huge problem," said Meghan Quint, one of the original team members of Opala and current SVP of Product and Customer Success.

When Premiera spun out Opala as an independent company, the team debated whether to build or buy. The answer quickly became clear. Building a clinical-grade identity layer in-house—one that could scale across multiple payers and health systems, configure quickly for new clients, and maintain accuracy as data volumes grew—would have consumed the engineering capacity they needed to build the platform they actually wanted to bring to market. Identity was too important to get approximately right, and too complex to build well while building everything else.

So they made the infrastructure decision early. They partnered with Verato to handle identity resolution, and focused their own engineering on what differentiated Opala in the market.

Opala's platform—HealthSynq™—ingests and normalizes claims, clinical, lab, and ADT data feeds in near real time, turning fragmented data from disparate sources into a complete, actionable picture of patient health. That is a technically demanding proposition. And it depends entirely on one foundational capability: knowing, with confidence, who every record belongs to.

Trusted identity has become part of how Opala positions itself in the market: not just an internal capability, but a signal of platform reliability. As Ken Chandler, CEO of Opala, put it:



“Verato has set the standard for trusted identity resolution, an essential foundation for interoperability. By ensuring patient data is accurately matched across systems, Verato enables Opala to build complete, longitudinal patient records and exchange them confidently at scale for our customers.”

Ken Chandler, CEO, Opala

Opala's story is instructive not because it is unusual, but because it reflects a decision increasingly made by health tech companies. The team knew the identity problem better than most, and still concluded that building it themselves was the wrong call. When the people most qualified to build something in-house decide to treat it as infrastructure instead, that is a signal worth paying attention to.

Spotlight: Opala Makes Interoperability Radically Easier

Opala is a data automation partner that makes payer-provider collaboration easier by giving payers, providers, and digital health innovators a complete, accurate, and timely picture of patient health.



For payers:

Opala simplifies compliance with interoperability mandates, reduces the operational burden of data exchange, and delivers the complete member views needed to support value-based care, risk adjustment, and population health programs.



For providers:

Opala closes gaps in patient records, improves care coordination across settings, and reduces the administrative friction that comes with fragmented data and manual reconciliation.



For health tech companies:

Opala provides a trusted data foundation that accelerates time to market, reduces integration complexity, and enables the kind of scalable, reliable data exchange that modern digital health platforms depend on.

Opala's mission is to be the most trusted source of patient insights and radically improve healthcare by making the data that powers it complete, accurate and actionable.

What buyers should look for:

Criteria for the next generation of health tech infrastructure

Not all identity solutions are built for what modern health tech platforms actually need. Many were designed for a different era, when data volumes were smaller, system landscapes were simpler, and interoperability was an aspiration rather than a mandate. Evaluating identity infrastructure with the same rigor applied to any other foundational technology investment means asking the right questions before committing to an approach.

These are the criteria that matter.



Cloud-native and built to scale

Identity infrastructure that was architected for on-premises deployment or that requires significant infrastructure overhead to operate does not belong in a modern health tech stack. The right solution is cloud-native by design and should scale with the platform without requiring architectural rework every time data volumes increase or new client environments are added..



Works across systems, products, and organizations

Identity infrastructure that only works within a single system or a single product line is not infrastructure, it's a point solution. The right foundation resolves identity consistently across EHRs, CRMs, claims systems, data platforms, and digital products. It works across organizational boundaries, supporting data exchange with partners, payers, providers, and networks without requiring a separate identity reconciliation effort for each new connection.



Purpose-built for healthcare complexity

Generic MDM solutions and open-source matching tools were not designed for the specific complexity of healthcare data. Provider directories, payer-provider relationships, clinical records, claims data, and consumer data all carry different characteristics, different error patterns, and different accuracy requirements. Healthcare identity requires healthcare-specific logic, which should come pre-built, not configured from scratch by the buyer's own engineering team.



Supports AI and analytics use cases

Identity infrastructure intended to power AI initiatives should be evaluated on its ability to produce clean, complete, and deduplicated populations, not just its ability to match records within a single system. Ask specifically how the solution handles duplicate resolution at scale, how it manages identity across longitudinal data spanning multiple years and multiple sources, and whether it can deliver the high-fidelity input data that AI models require to produce outputs worth trusting.



Reduces engineering lift and manual stewardship burden

The right solution should reduce operational burden over time, not increase it. Look for solutions that automate the majority of identity resolution algorithmically, reserve manual stewardship for genuine edge cases, and require minimal per-client or per-source configuration when the platform expands.



Accelerates time to value

A solution that requires six to nine months of configuration before it can be deployed for a new client creates a ceiling on how fast the platform can grow. Modern identity infrastructure should be pre-configured for healthcare data from day one, meaning a new data source or a new client environment can be onboarded in weeks, not months.



Supports interoperability and network-based exchange

Interoperability requirements are expanding in scope and complexity. Identity infrastructure that was designed for internal use only will struggle to serve as the foundation for network-based exchange models—TEFCA, Da Vinci, FHIR-based APIs, and ecosystem connectivity frameworks that require identity to be resolved consistently across organizational boundaries. Evaluate whether the solution was built with network-scale interoperability in mind.

These criteria won't all carry equal weight for every organization. But they share a common purpose: helping leadership teams distinguish between identity solutions that solve today's problem and identity infrastructure that supports the platform's coming years of growth.

Conclusion

The next modernization move is identity

Health tech companies have already made the case for modern infrastructure, repeatedly and convincingly. They moved to cloud, adopted purpose-built analytics platforms, and invested in modern CRMs and AI tooling because the logic was clear: commodity plumbing should not drain strategic resources.

That same logic applies to identity. And yet for many organizations, identity has been left behind, still handled through homegrown logic, legacy MDM platforms, or point solutions never designed to serve as enterprise infrastructure.

That gap is no longer a background inefficiency. It is an active constraint on product velocity, AI reliability, interoperability, and growth. The companies that install identity as infrastructure will launch faster, integrate more easily, support AI more confidently, and deliver more reliable experiences at scale. The ones that continue treating it as a side project will feel that constraint compound.

The modernization work is not finished. For most health tech companies, identity is the next move.

Ready to rethink identity's role in your tech stack?

Explore how Verato can help you build a trusted identity foundation.

Book a demo

Take the identity maturity assessment

Learn more at verato.com

Sources

¹ Pew Charitable Trusts, Enhanced Patient Matching Is Critical to Achieving Full Promise of Digital Health Records, October 2018. https://www.pew.org/-/media/assets/2018/09/healthit_enhancedpatientmatching_report_final.pdf

² RBC Capital Markets, The Healthcare Data Explosion. IDC/Seagate DataAge research cited therein. Healthcare data CAGR of 36% projected through 2025. http://www.rbccm.com/en/gib/healthcare/episode/the_healthcare_data_explosion

³ McKinsey & Company, Generative AI in Healthcare: Current Trends and Future Outlook, Q4 2025 survey of 150 U.S. healthcare leaders across payers, health systems, and healthcare services and technology groups. <https://www.mckinsey.com/industries/healthcare/our-insights/generative-ai-in-healthcare-current-trends-and-future-outlook-2025>

⁴ Capstone Partners, Healthcare IT M&A Update, January 2025. <http://www.capstonepartners.com/insights/article-healthcare-it-ma-update>

⁵ American Hospital Association, Fact Sheet: Telehealth, February 2025. Kaiser Family Foundation data cited therein. <http://www.aha.org/fact-sheets/2025-02-07-fact-sheet-telehealth>



Verato®, the identity intelligence company, solves the problem that drives everything else—Knowing Who Is Who™. Proven in healthcare. Trusted everywhere. Verato MDM Cloud™ delivers unprecedented identity intelligence by combining industry-leading identity accuracy with enterprise-wide data mastering to unify, enrich, and manage identities across systems of insight, engagement, and record—unlocking more accurate, actionable data insights. With complete, accurate, and trusted identity intelligence, Verato improves business results by powering exceptional consumer experiences and patient outcomes, fueling business growth, accelerating AI readiness, boosting productivity, and reducing costs across every business process that consumes identity. For more information, visit verato.com.