Hylaine

Building & Unifying Data in the Cloud:

An execution playbook for insurance leaders

Execution is where data strategies succeed or stall.

This paper provides a practical framework for insurance organizations working to build, unify, and scale data in the cloud. Drawing on real-world engagements, it focuses on how the work actually gets done.

Justin Goff & Ryan McElroy

Table of Contents

Anticipated challenges: What's around the corner

Challenges shared across all insurance

- Data Fragmentation
- Compliance Imperative
- Legacy System Burden
- AI Readiness Gap
- Governance Requirements

Where insurance types diverge

Healthcare / MCO

CMS mandates, HIPAA, real-time eligibility

P&C

State-by-state filing, catastrophe modeling, IoT

Life & Annuity

Multi-decade policy data, actuarial complexity

Specialty / Surplus

Emerging risk models, flexible data schemes

Reinsurance

Treaty aggregation, quarterly cadence, cross-carrier

© 2026 Hylaine

Introduction: Execution is the Real Challenge

Insurance organizations don't struggle with strategy. They struggle with execution. Most have already committed to cloud, data modernization, and AI. But once the work begins, progress slows because legacy systems remain in place, data stays fragmented, and teams struggle to align on how to move forward in a way that holds up at scale.

You can see the pressure building. AI initiatives depend on data that is consistent and accessible. Regulatory expectations continue to rise. Investment in technology is accelerating across the industry (see graphic below). Most organizations are moving forward, but many are doing so without a clear path to unify data in a way that supports day-to-day operations.

This paper focuses on execution, outlining how to build, unify, and operationalize data in the cloud using approaches that work in complex environments. You'll see where these efforts tend to break down, what strong governance and data reliability look like in practice, and how to move from early progress to sustained adoption in your own organization.

Investment in technology is accelerating across the industry

$173B

U.S. Insurance IT spend projected for 2026

Source

91%

of banks & insurers actively migrating to cloud

Source

$13.9B

global AI-in-insurance market projected 2026

Source

95%

of enterprise AI pilots fail to deliver ROI

Source

Strategy gets you to the starting line.

Execution gets you across the finish.

© 2026 Hylaine

Where Data Transformations Break Down

Insurance organizations operate in very different environments. A Medicaid managed care plan processing eligibility in real time does not look like a life insurer managing decades of policy data, or a specialty carrier pricing cyber risk. The details vary. The underlying challenges do not.

Across the industry, some issues show up consistently:

1. **Data fragmentation –** Data is spread across policy, claims, billing, and third-party systems that were never designed to work together.
2. **Compliance imperative –** Data must remain accurate, traceable, and secure under continuous regulatory scrutiny.
3. **Legacy system burden –** Outdated systems continue to shape how data is structured and interpreted over time.
4. **AI readiness gap –** Data lacks the quality, structure, and governance required to support trusted AI and analytics at scale.
5. **Governance requirements –** Data must meet strict standards for lineage, security, and auditability across distributed environments.

Where insurers diverge is in how these pressures show up. Some operate in highly dynamic regulatory environments, others in more stable but fragmented ones. Some depend on real-time data flows, others on periodic processing. Product complexity can either demand flexible data models or make them difficult to change. These differences matter, but they are not what causes most data initiatives to stall.

The breakdown happens in how data is interpreted and carried forward.

Cloud efforts are often treated as lift-and-shift migration programs. Move the data. Reconnect the systems. But legacy data is not just stored differently—it reflects years of embedded business logic and regulatory requirements. From on-prem applications, undocumented stored procedures, and the institutional knowledge of people approaching retirement. A field is not just a field. It carries meaning shaped by the system and process that produced it.

In practice: Fortune 25 managed care organization

Insurance claims data existed across multiple on-prem and SaaS systems with overlapping and conflicting records. A lift-and-shift approach would have broken compliance requirements. Data had to be restructured in the cloud to meet new application needs while preserving its original meaning, auditability, and regulatory integrity.

Follow this client's progress as they build and unify data in the cloud throughout this paper in the "In Practice" sections.

© 2026 Hylaine

Unifying data requires reconciling that meaning. It requires preserving compliance controls that may not transfer automatically. It requires maintaining transactional integrity so that records remain complete, traceable, and audit-ready throughout the transition.

This is where many efforts slow down or fail. Not because the technology cannot support the move, but because the work is treated as a technical exercise instead of an operational one.

The real challenge: More than moving data

What People Assume	What it Actually Requires
⚠ Simple lift-and-shift migration	✓ Reconciling meaning across systems
⚠ One-to-one data mapping	✓ Preserving embedded business logic
⚠ Technology problem only	✓ Business transformation & technology
⚠ Point-in-time project	✓ Ongoing program with sustained adoption

In practice: The compliance dimension

Insurance claims data existed across multiple on-prem and SaaS systems with overlapping and conflicting records. A lift-and-shift approach would not have satisfied both the new application's requirements and existing compliance and auditing regulations. Data had to be restructured in the cloud to meet new application needs while preserving its original meaning, auditability, and regulatory integrity.

© 2026 Hylaine

Governance and Reliability: The Foundation for Execution

If data unification breaks down when meaning is not preserved, governance is what keeps that meaning intact as the work moves forward.

Before any large-scale cloud effort gains traction, your teams need a shared understanding of their data. What does it represent? Who owns it? How should it be used? And what standards must it meet? Without that alignment, inconsistencies multiply quickly across systems, teams, and use cases.

In insurance, this carries additional weight. The data being unified today will support not only reporting and operations, but also future AI models and decisioning. That raises the bar. Data must be accurate and complete, but also explainable, traceable, and monitored for bias over time. Issues that sit quietly in legacy systems tend to surface when data is reused in new ways. For instance, errors in data movement can result in millions of dollars of compliance liabilities for insurance companies.

Data reliability engineering (DRE) is what makes this practical. DRE encompasses observability, data quality monitoring, lineage tracking, anomaly detection, testing, and root cause analysis. Governance defines the expectations, and DRE enforces them through continuous validation.

In a cloud environment, that means monitoring how data changes as it moves between systems. It means tracking lineage, testing transformations, and detecting anomalies early. Most importantly, it confirms that even when data is restructured, it still reflects the same business meaning and meets the same compliance requirements.

How data changes as it moves between systems

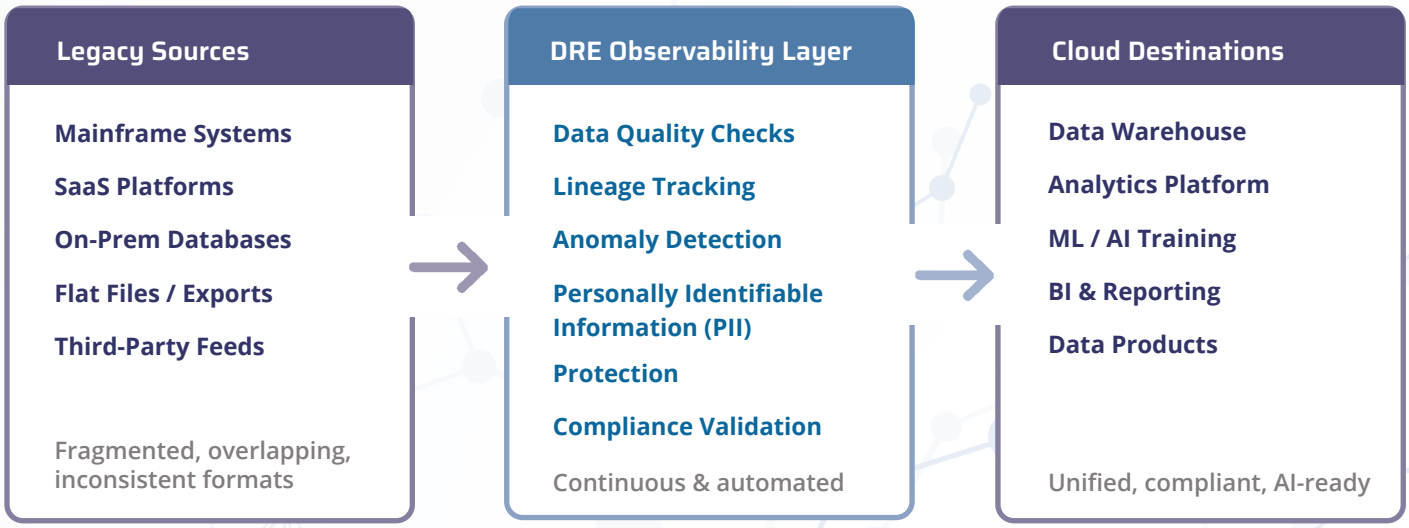

© 2026 Hylaine

Manual oversight does not scale here. The volume of data, the number of dependencies, and the regulatory stakes require automated, ongoing validation.

A governance council helps maintain alignment as this work progresses. Representatives from IT, compliance, actuarial, claims, underwriting, and executive leadership ensure that decisions reflect both operational needs and regulatory realities. This is not about adding process. It is about preventing drift as complexity increases.

Data governance and AI governance should also be treated as related, but distinct. One establishes a reliable foundation. The other governs how that foundation is used.

Data vs. AI governance: Related, but distinct

Data Governance	AI Governance
Define data ownership & stewardship	Bias detection & demographic drift monitoring
Set quality standards & SLAs	Model explainability requirements
Establish lineage requirements	Regulatory transparency standards
PII protection & compliance automation	Responsible AI policies & guardrails
Governance council with cross-functional reps	Separate discipline, tightly connected to data

In practice: Validating data in motion

For this client, DRE was implemented by Hylaine as an observability layer between legacy systems and cloud platforms. Automated checks confirmed that data met quality, compliance, and audit requirements as it moved—preserving PII protections, lineage, and regulatory alignment even when the structure changed.

© 2026 Hylaine

Execution Framework: Moving from Alignment to Scale

With governance and reliability in place, execution becomes a matter of sequence. The work typically unfolds across three phases, where each phase builds on the one before it. When teams skip steps or try to accelerate one phase in isolation, progress slows later.

This is not a linear handoff. The phases overlap. But the order still matters.

A phased approach for execution

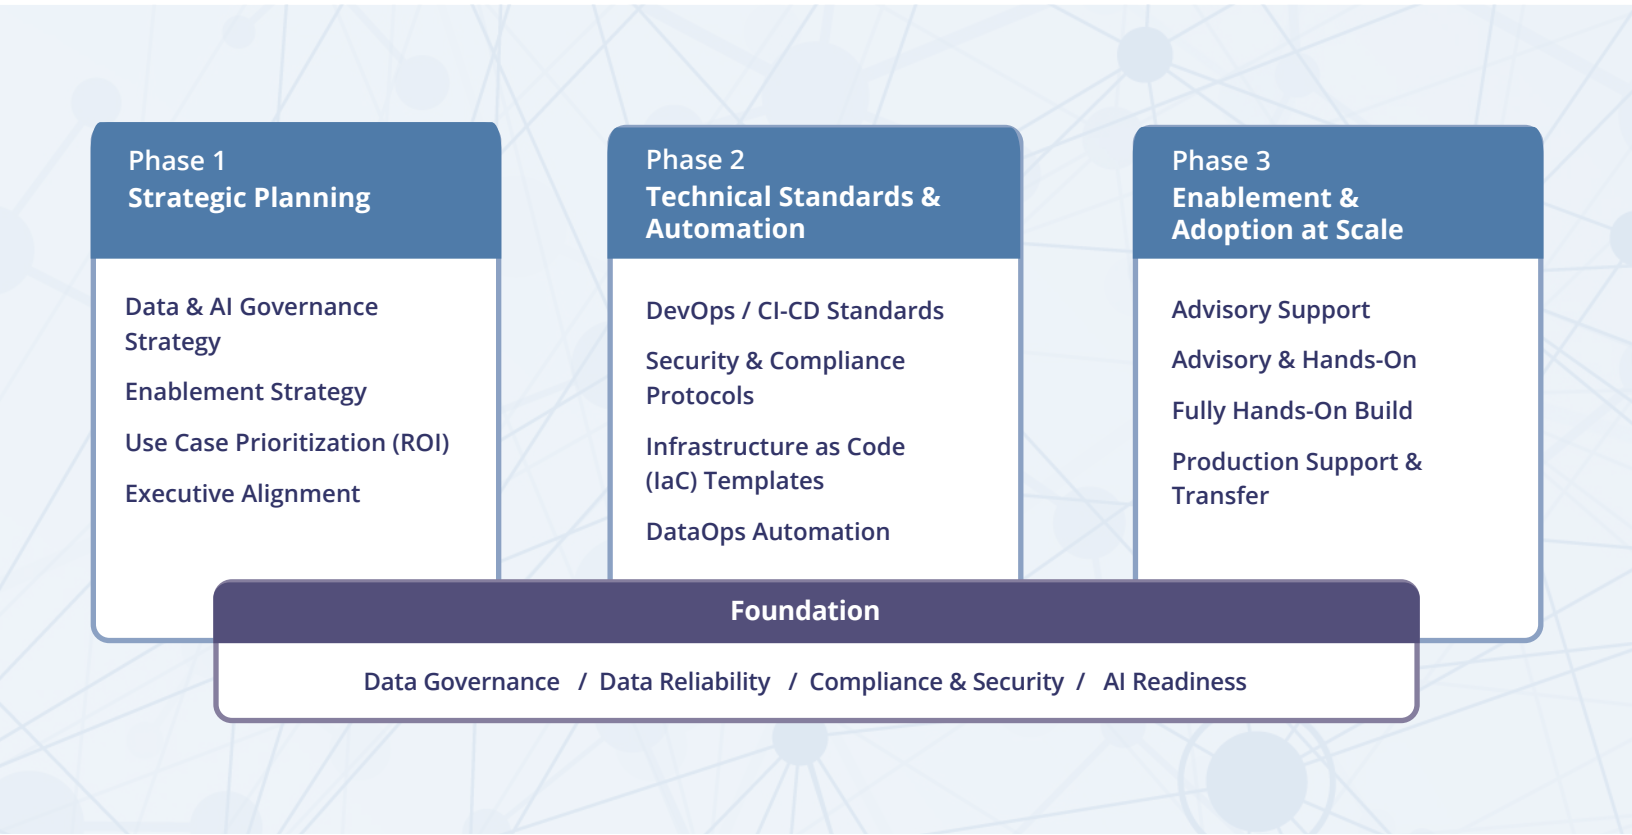

© 2026 Hylaine

Phase 1: Strategic planning

Execution starts with a focused form of strategy. Not an open-ended assessment, but a defined effort that produces decisions teams can act on quickly.

The first priority is governance alignment. Data ownership, quality standards, and usage expectations need to be clearly defined and documented. This includes how data will support future AI use cases, not just current reporting or operations.

From there, the work shifts to enablement planning. This is where strategy becomes practical. Which use cases come first? What outcomes do they need to deliver? And how will progress be measured?

A short-term execution plan, typically over 90 days, creates initial momentum. A longer-term roadmap ensures that early efforts connect to broader priorities.

The key is focus. When use cases are tied to measurable business impact, your teams can demonstrate progress early. That visibility is what sustains executive support and keeps the program moving.

Data & AI Governance Strategy

Establish governance policies, define data ownership, set quality standards, and create the framework for AI readiness. The output is a documented strategy that includes both technical implementation standards and organizational policy.

Enablement Strategy

Transfer governance into an execution playbook, identify highest-priority use cases grounded in measurable business impact and ROI. Include a 90-day sprint plan and a 12-month roadmap. Grounding in business value is essential for sustained executive buy-in.

© 2026 Hylaine

Phase 2: Technical standards and automation

Once priorities are clear, execution moves into how the work gets done.

Many organizations struggle here. Without clear standards, cloud environments quickly become inconsistent. Different teams adopt different patterns. Over time, the new environment starts to resemble the one it replaced.

The first step in Phase 2 is defining technical standards. These include infrastructure patterns, security requirements, identity and access controls, CI / CD processes, and data platform conventions. In insurance, security and compliance are central, and every deployment needs to meet the same baseline expectations.

But standards alone are not enough. Teams need a way to apply them consistently without slowing down.

That is where automation comes in.

Infrastructure as code templates translate standards into something usable. When a team provisions a new data pipeline or application, the required configurations are already in place. Security policies, tagging, and monitoring are built in. The team does not need to interpret the standard. They inherit it.

This approach changes how teams work. Instead of designing from scratch, they start from a known baseline. That reduces friction early and prevents inconsistency later.

It also improves long-term manageability. When environments follow consistent patterns, they are easier to support, audit, and evolve.

Establish Technical Standards

- DevOps & CI / CD pipeline standards
- Identity & access management (IAM)
- IaC conventions
- Security assessment protocols
- Cost optimization & resource tagging
- Data platform conventions (OLTP / OLAP)

Automate Through Templates

- Pre-configured IaC with tagging built in
- Security policies inherited by default
- Templated database & compute deployments
- Containerized Kubernetes packages
- Monitoring dashboards auto-provisioned
- Plug-and-play deployment patterns

In practice: Building standards into deployment

For this client, Hylaine created templates that were used to preconfigure security, access controls, cost tagging, and monitoring. Teams could deploy new workloads with these elements already in place, reducing variation and avoiding rework.

© 2026 Hylaine

Phase 3: Enablement and adoption at scale

The final phase is where the framework is tested. Individual teams begin adopting the platform, migrating workloads, and working within the established standards.

However, this does not happen evenly. Different teams bring different levels of technical maturity. Typically, some move quickly, while others need more support.

A dedicated enablement team helps bridge that gap.

In practice, engagement tends to take three forms: advisory for technically mature teams, advisory plus limited hands-on for teams with gaps in specific areas, and fully hands-on for teams that need the enablement team to lead the build.

The structure can vary. What matters is consistency. Every team works from the same templates, standards, and governance model.

As adoption grows, the focus shifts to sustainability. Teams need to operate what they build. That requires training, documentation, and a clear transition of ownership.

Not every team will manage production independently. In those cases, a managed support model may be appropriate. The important point is clarity. Each workload should have a defined owner and a defined support approach.

At this stage, data and infrastructure are no longer separate efforts. Scaling data depends on scaling the platform that supports it. The enablement model keeps both aligned as adoption expands.

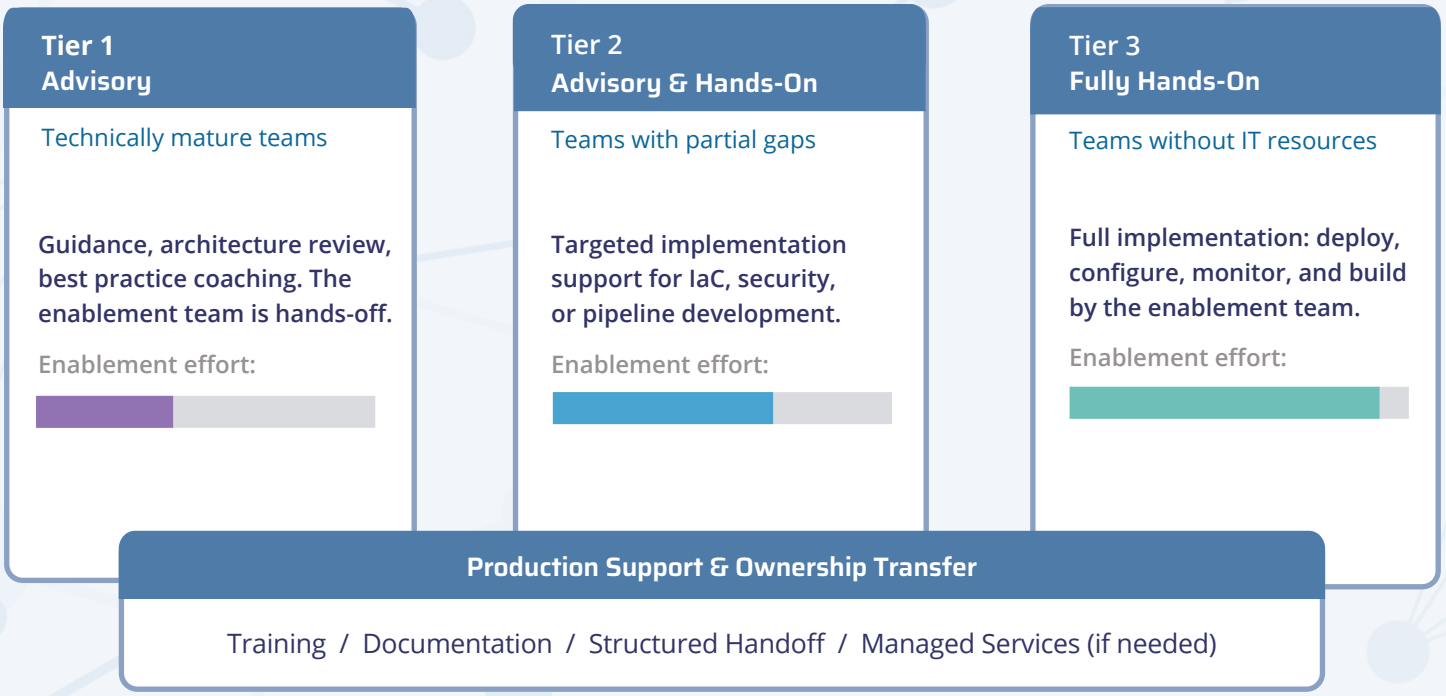

© 2026 Hylaine

Unification Without Consolidation: The Federated Approach

Data unification is often interpreted as consolidation. Move everything into one platform. Standardize access in a single location.

In insurance, that approach breaks down quickly.

Regulatory requirements, security constraints, and organizational structure often require data to remain in separate environments. For many organizations, forcing consolidation introduces more risk than it removes.

A federated approach takes a different path. Data stays where it is. What changes is how it is accessed and used.

A real-world architecture with a federated approach

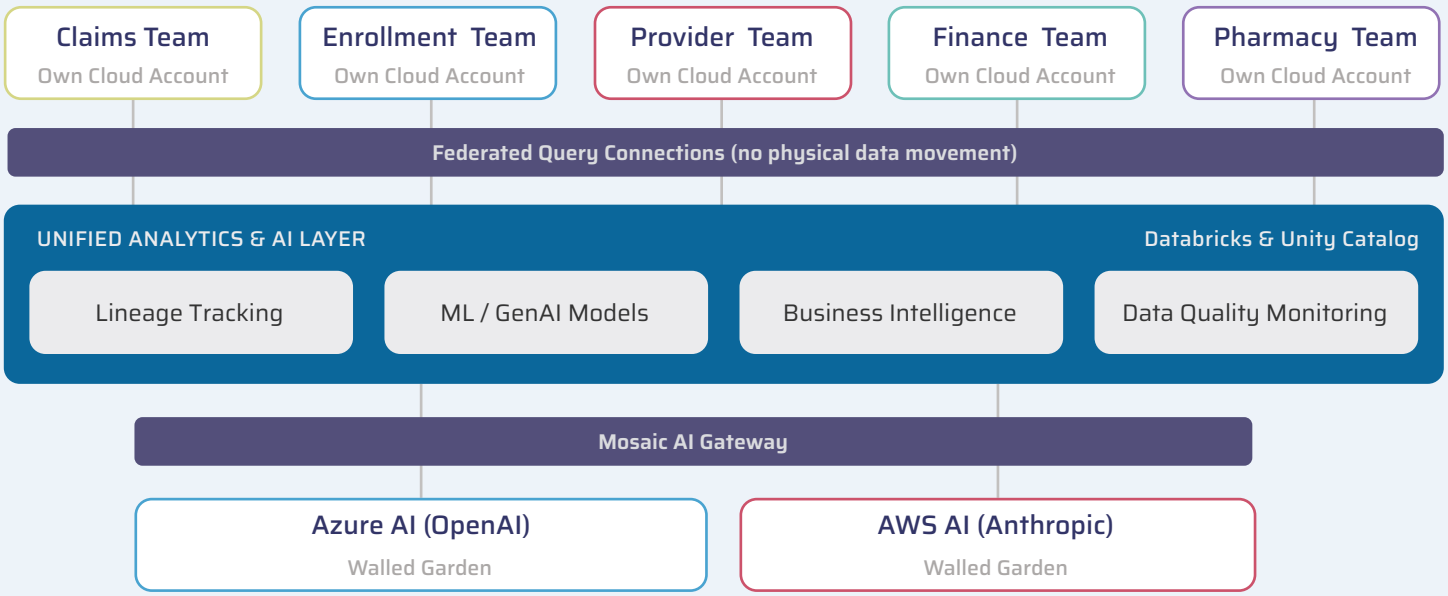

A unified analytics and AI layer allows teams to query, join, and analyze data across systems without physically relocating it. This preserves existing security and compliance boundaries while still enabling broader visibility. Teams maintain control over their data. The organization gains the ability to work across it.

© 2026 Hylaine

In practice: Federated architecture at scale

For this client, Hylaine helped to implement Databricks as the centralized compute layer, with Unity Catalog providing lineage tracking across disparate sources. A Mosaic AI gateway connected Azure and AWS, enabling access to proprietary foundation models within a walled-off, security-assessed environment. This enabled individual teams to maintain control of their own systems. Federated queries enabled analytics, machine learning, and AI workloads to operate across these sources, with lineage tracked end-to-end. The result was shared visibility without requiring data consolidation.

This model is particularly effective for organizations that operate across multiple lines of business, manage sensitive data, or have grown through acquisition. It reflects how most insurance environments actually function.

Federation is not a workaround. It is a deliberate design choice that supports scale while respecting constraints.

Tools like Databricks, Collibra, Alation, and Manta can extend this approach by adding visibility into data lineage, governance, and usage across distributed environments.

© 2026 Hylaine

Getting Started: From Playbook to Action

Execution tends to stall when the starting point is unclear.

Most organizations benefit from a focused assessment across four areas.
This checklist can help start that assessment:

Data governance maturity

- ☐ Do you have a documented data governance maturity strategy?
- ☐ Have you established a Governance Council?
- ☐ Are quality standards enforced?

Infrastructure readiness

- ☐ Have you implemented standardized IaC templates?
- ☐ Do you have consistent deployment patterns across teams?

Enablement capacity

- ☐ Are there internal resources in place for team adoption?
- ☐ Is external support required?

Compliance posture

- ☐ Do you have an auditable process?
- ☐ Is there a source-to-destination lineage?
- ☐ Has automated PII protection been established?

This does not need to be a long process. A targeted diagnostic can establish a clear baseline in a matter of weeks.

From there, your progress depends on sequencing. Identify the highest-value use cases. Build a phased roadmap. Put the enablement model in place early so adoption can scale.

The work itself is often straightforward. Building pipelines. Standardizing data. Validating outputs. It's not new. But the expectations around speed, scale, and accountability have changed.

Organizations that move effectively tend to combine internal expertise with external support. Your team brings context—how the business operates, where the risks sit, what matters most. External specialists bring pattern recognition and execution experience across similar environments. Together, the work moves faster and with fewer false starts.

The playbook is established. The next step is applying it to your environment, with clarity on where to begin and how to sustain progress.

Start with a clear view of your current state. Then act where it counts—and create the momentum that carries this work from isolated wins to organization-wide impact.

© 2026 Hylaine

About Hylaine

Hylaine, a values-first technology consulting firm, works with property & casualty, life & annuity, health, specialty & Lloyd's, and reinsurance to modernize and simplify the technology behind the insurance business. From outdated policy and claims systems to fragmented data and regulatory demands, Hylaine helps insurers accelerate adoption of digital-first operating models, reduce risk, and improve speed to market. Its capabilities span strategy, systems integration, process optimization, and enterprise technology modernization.

What sets Hylaine apart is its commitment to a values-first approach that stands for partnerships over transactions, doing what's right over what's easy, honesty without exception—no bait and switch—ever, and transparency in everything. This approach has helped Hylaine earn 4 straight appearances on the Inc. 5000 list—including the 2025 ranking—thanks to a reported 216% 3-year revenue growth.

Authors

Justin Goff, Director of Technical Delivery

With nearly two decades of experience in advanced analytics and AI-driven solutions, Justin focuses on delivering scalable, real-world outcomes—helping organizations turn complex data challenges into operational results that drive performance and growth.

Ryan McElroy, Vice President, Technology

With deep experience in data governance, AI readiness, and data architecture, Ryan helps clients realize the real value of their data, not just for today's reporting needs, but for tomorrow's innovations.

Connect with us and see what better looks like.

hylaine.com
info@hylaine.com
Schedule a Complimentary Discovery Session