

eBook

Cloud Modernization and Tools Consolidation: A Technical Leader's Guide to Scalable, Secure Growth

Where data
& AI come to 

Contents

The Data Foundation Crisis Hiding in Plain Sight	3
Why Now: The Forces Making Legacy Data Infrastructure Untenable	4
The Hidden Cost of Tool Sprawl	6
The Cost of Inaction	7
What a Modern Cloud Data Foundation Actually Looks Like	9
6 Core Capabilities of a Unified Data Platform	10
Navigating the Transformation: Why Cloud Modernization Efforts Stall and How to Avoid It	11
Charting Your Optimal Course to Modernization	12
Taking Stock: Are You Ready to Modernize?	13
From Reactive to Future-Ready	14
Next Steps	15
About Us	16

The Data Foundation Crisis Hiding in Plain Sight

It's undeniable: the age of AI is upon us. Generative artificial intelligence (GenAI) adoption jumped to 69% in 2025, and agentic AI, though emerging, has already been adopted by nearly half of companies.¹ As with GenAI, the success of agentic AI hinges on a key factor: your data. Half of CDOs still cite data quality and retrieval concerns as the most common challenge in migrating GenAI and agents into production.²

For enterprises racing to deploy AI and agentic workloads, this represents a harsh reality of stalled roadmaps, missed milestones and proof-of-concepts that never make it to production. But the inability to scale agentic AI projects rarely stems from a lack of vision, resourcing or skills. Instead, the failure point is often a lack of trusted context supported by a solid data infrastructure.

Years of “best-of-breed” tool selection decisions, which were rational in isolation, have compounded into fragmented data landscapes. These legacy environments were never designed for the modern mandate: real-time analytics, AI at scale and agentic workloads that demand clean, governed and trusted data on demand.

It's no wonder 85% of data leaders say modernization is critical to enable AI/ML.³ Future-ready organizations aren't waiting for a perfect migration window; they are actively consolidating fragmented ecosystems onto unified, cloud-native foundations. By building a robust data infrastructure, these leaders are providing the trusted context necessary to make AI ambitions executable at scale.

The mandate for technology leaders is clear: modernize and consolidate data infrastructure now, or watch AI innovation stall at the pilot stage.

This guide explores why fragmented legacy environments are fundamentally incompatible with the AI-ready enterprise. We will break down the architecture of a future-proof cloud data foundation and how forward-thinking leaders are navigating this shift safely and strategically.

¹ Informatica, CDO Insights 2026.

² ibid

³ Informatica, The State of PowerCenter Modernization 2025.

Why Now: The Forces Making Legacy Data Infrastructure Untenable

The AI Ambition Gap

Across industries, the directive from the boardroom is unmistakable: AI is a strategic priority. In fact, 99% of executive leadership states that investments in data and AI are a top organizational priority.⁴ Budgets are allocated, and leadership is aligned. However, while the pressure to deliver is intensifying, a critical gap has emerged: corporate ambition is outpacing technical infrastructure.

On the ground, technical teams are inheriting the consequences of data environments that were never designed for AI-scale demands. Legacy pipelines that move data in batch cycles cannot support the real-time data access that AI models require. **Governance** frameworks bolted onto aging architectures cannot enforce the trust and **lineage** that production AI demands. Furthermore, the rise of agentic AI — autonomous systems designed to handle intricate, multi-layered workflows — imposes a significant burden on data architectures that were originally engineered for a different era.

This leaves technical teams on the ground to grapple with data environments that were never designed for the rigors of AI. This disconnect manifests in three primary ways:

- **Latency vs. real-time demands:** Legacy pipelines built for batch cycles are fundamentally incapable of supporting the real-time data access modern AI models require.
- **Static governance:** Governance frameworks bolted onto aging architectures cannot enforce the rigorous trust and lineage tracking required for production-grade AI.
- **The agentic strain:** The shift toward agentic AI systems, which act autonomously to complete complex, multi-step tasks, places unprecedented strain on infrastructure built for a fundamentally different era.

The result is a widening gap between leadership's expectations and the actual capabilities of the data layer. While 93% of companies say that interest in AI has led to a greater focus on data,⁵ the execution remains flawed: 63% percent of organizations training AI models have wasted resources doing so on bad data.⁶

⁴ NewVantage Partners, 2026 AI & Data Leadership Executive Benchmark Survey.

⁵ ibid

⁶ Salesforce, State of Data and Analytics, 2025.

Why Now: The Forces Making Legacy Data Infrastructure Untenable (continued)

The consequences of this disconnect are stark. According to MIT, despite \$30–40 billion in enterprise investment into GenAI, 95% of organizations are getting zero return.⁷ Technical hurdles further compound these failures. Specifically, only 40% of companies surveyed by Deloitte feel that their data management is highly prepared to support AI adoption.⁸

These figures translate into a familiar, high-stakes reality: pipelines breaking under new loads, data quality issues surfacing mid-deployment and governance gaps creating compliance exposure at the worst possible moment. When these failures occur, a lack of traceability often delays critical remedial action.

For technical leaders, closing this gap is no longer a future consideration. It is the immediate operational challenge, and legacy infrastructure remains the primary obstacle standing in the way.



⁷ MIT, The GenAI Divide: State of AI in Business 2025.

⁸ Deloitte, The State of AI in the Enterprise, 2026.

The Hidden Cost of Tool Sprawl

This fragmentation was not an overnight occurrence. The **extract, transform, load (ETL) platform** was implemented years ago for batch processing, data quality tools arrived as compliance tightened and the integration layer was bolted on alongside the first cloud applications.

While each decision was rational in isolation, they have collectively evolved into a fractured architecture that no single team fully understands — or truly owns.

This is the hidden weight of tool sprawl. For most enterprises, the burden is far heavier than vendor licensing fees suggest. The true costs are operational:

- **Brittle infrastructure:** Point integrations and pipelines require constant maintenance and frequently break under modern data volumes.
- **Fragmented logic:** Duplicated transformation logic is scattered across platforms, leaving the organization with no single source of truth.
- **Budget inefficiency:** Redundant ETL/ELT tools drive up costs and complexity, while disconnected platforms further fragment financial resources.

- **Unreliable standards:** Inconsistent data quality standards vary across systems, making it nearly impossible to build governed AI and analytics pipelines with confidence.
- **Compliance risk:** Security and access control policies are enforced differently across tools, creating gaps that auditors and regulators will eventually uncover.

Critically, this is not a staffing problem. Hiring more engineers to maintain a fundamentally fragmented architecture does not resolve the underlying structural issue. It scales the complexity alongside the headcount, creating dependency on individuals who carry institutional knowledge that exists nowhere in documentation.

When you recognize tool sprawl as an architectural liability rather than a management challenge, you are better positioned to make modernization decisions that actually work.

The Cost of Inaction

Over the next 18 to 24 months, enterprises that delay modernization will face a convergence of pressures making the transition increasingly difficult and expensive.

AI-driven competitors are gaining ground — not necessarily through superior models, but because their data infrastructure allows them to iterate, test and deploy at speeds that legacy environments structurally cannot match.

The talent implications are equally significant. Experienced data engineers are gravitating toward modern, cloud-native environments, making it difficult to retain them on aging on-premises platforms.

The regulatory dimension adds further urgency as global data privacy mandates, AI governance frameworks and ESG reporting requirements tighten across every major market. Enterprises operating fragmented, manually governed data environments are accumulating compliance exposure with every passing quarter.

Industry data reinforces these stakes. BCG research shows a widening gap between AI leaders and laggards, with the difference increasingly attributable to data infrastructure readiness rather than strategy.⁹ The window to close that gap competitively is narrowing.

⁹ BCG, The AI Adoption Puzzle: Why Usage Is Up But Impact Is Not, 2025.



The Cost of Inaction (continued)

How Data Infrastructure Challenges Impact Business Outcomes for Modern Enterprises

IT Challenge	Business Impact
Migration of large volumes of data from on-premises systems to the cloud	Potential system downtime, high migration costs and data integrity issues
Maintenance-intensive on-premises system	Elevated costs of maintaining expensive hardware, infrastructure and data centers
Difficulty integrating existing on-premises systems with new cloud services	Higher likelihood of system disruption plus poor interoperability between legacy systems and modern cloud technologies
Overly complex tasks for managing and updating software	Excessive time and resources are required to keep software and systems current
Threats to the security of sensitive data	Increased effort required to protect against potential vulnerabilities and comply with industry regulations
Limited ability to deploy new services and applications due to legacy infrastructure	Reduced business agility and innovation
Fragmented data spread across multiple tools, platforms and vendors	Higher total cost of ownership, duplicated capabilities and inconsistent data across the enterprise
Managing multiple integration, quality and governance tools post-migration	Continued operational complexity, higher licensing and support costs and reduced ROI from cloud investments
Difficulty standardizing security and compliance controls across fragmented platforms	Increased risk exposure, inconsistent policy enforcement and audit complexity

What a Modern Cloud Data Foundation Actually Looks Like

An Architecture That Unlocks AI Readiness

The shift from fragmented data infrastructure to a unified, cloud-native foundation is more than a technological upgrade; it is an architectural shift that redefines how data is ingested, governed, trusted and delivered across the enterprise.

In fragmented environments, governance is reactive. Data quality checks are layered onto pipelines never designed for them, lineage is manually reconstructed and compliance is enforced inconsistently across disparate platforms. This results in data that arrives late, contains errors or lacks the context required for confident action.

A **modern cloud data foundation** inverts the legacy model by building governance into the architecture from the ground up rather than bolting it on after deployment. Every data asset carries metadata from the moment it enters the ecosystem. Lineage is tracked automatically. Quality rules are enforced uniformly across all sources and domains. This approach makes AI readiness possible at scale.

The other critical shift is the role of metadata intelligence. In a unified platform, metadata automates the work of entire teams of data stewards — automatically discovering assets, mapping relationships, flagging anomalies and enforcing policies. This allows you to scale your data operations without proportionally increasing headcount. According to Gartner®, “by 2027, organizations that prioritize semantics in AI-ready data will increase their GenAI model accuracy by up to 80% and reduce costs by up to 60%.”¹⁰

¹⁰ Gartner Press Release, “Gartner Announces the Top Data & Analytics Predictions,” June 17, 2025. GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved.

6 Core Capabilities of a Unified Data Platform

Not all cloud data platforms are created equal. As you evaluate modernization options, the specific capabilities of your chosen platform will determine how quickly you can deliver AI-ready data and govern it at scale. Here's what to look for:

- 1. Comprehensive platform:** A genuinely unified data platform does more than replace individual tools with cloud-hosted equivalents. It converges data integration, governance, quality and cataloging into a single, coherent system where each capability reinforces the others. By consolidating fragmented tools into a single, AI-powered foundation, you can reduce complexity, cost and operational risk across the entire data lifecycle.
- 2. Cloud-native, elastic data integration at enterprise scale:** Unlike legacy ETL platforms that require manual scaling and fixed infrastructure, a truly cloud-native integration layer expands and contracts elastically with data volumes. It handles batch, streaming and real-time workloads across hybrid and multi-cloud environments without requiring pipelines to be re-engineered every time business demands change.
- 3. API and application integration for modern architectures:** Beyond the core data layer, modern architectures require robust API and application integration capabilities. Because enterprise resource planning (ERP) systems, SaaS platforms and operational applications act as both consumers and sources of data, a unified platform must connect to and orchestrate across them all without requiring separate integration tooling.

- 4. AI-powered metadata intelligence:** A unified platform should activate metadata rather than just store it. Embedded metadata intelligence automatically discovers data assets, maps relationships across the ecosystem, tracks end-to-end lineage and flags quality anomalies before they reach downstream consumers. For AI initiatives, this intelligence ensures models are trained on data that is understood, trusted and traceable back to its source.
- 5. End-to-end data governance and data quality:** Data assets cataloged in the system are immediately available for governance policy enforcement, with lineage flowing through the entire data lifecycle without manual reconstruction. Holistic data quality and observability ensure that data integrated through the platform is automatically subject to established quality rules.
- 6. Security and privacy posture built into the platform:** Security and compliance cannot be afterthoughts. A unified data platform should enforce data privacy policies, access controls and regulatory requirements uniformly through controls embedded in the platform itself, rather than through a separate compliance layer added on top.

When evaluating options, technical leaders should prioritize data platforms that deliver these capabilities under a single metadata layer. This ensures consistent governance and quality standards are applied regardless of where data originates or how it flows.

Navigating the Transformation: Why Cloud Modernization Efforts Stall and How to Avoid It

The decision to modernize to the cloud is rarely the difficult part. The challenge lies in executing a modernization initiative without destabilizing the data pipelines, workflows and data services your business currently depends on, while simultaneously building a future-proof architecture capable of any scale.

This is where most efforts falter. Organizations often underestimate the complexity of their existing workload inventory, leading to cascading disruptions when attempting to migrate everything at once. Conversely, taking an overly cautious, piecemeal approach can cause migrations to drag on for years, accumulating costs without delivering meaningful progress.

In either case, the fears are legitimate:

- Mission-critical data pipelines breaking mid-migration
- Losing months of productivity during transition
- Lacking internal skills to execute a streamlined modernization program
- Executive pressure to show results quickly while managing enormous hidden complexity

Enterprises that modernize successfully follow a different model.

First, you need to distinguish clearly between cloud migration and cloud modernization. Migration is the process of moving applications, data and workloads from on-premises or other environments to the cloud. Cloud modernization goes further by updating, consolidating and optimizing those assets to leverage cloud-native technologies for improved performance, scalability and agility.

If migration is about relocation via lift-and-shift approaches, modernization is about transforming and future-proofing enterprise data architecture.

Charting Your Optimal Course to Modernization

The path from legacy infrastructure to a unified, cloud-native data foundation is complex. Even with a strategic modernization and consolidation approach, you must address challenges such as:

- re-architecting legacy systems for cloud compatibility
- managing hybrid and multi-cloud complexity
- ensuring data security, governance and regulatory compliance in a distributed and consolidated cloud environment

Overcoming these obstacles requires a coordinated approach that aligns with strategic objectives, accelerates the modernization journey, maximizes asset reuse to protect investments and maintains operational continuity throughout.

For organizations running on-premises data tools, whether legacy third-party platforms or Informatica products like PowerCenter, a **structured cloud modernization and consolidation program** can dramatically reduce your migration risk, cost and timelines. It transforms what would otherwise be a multi-year disruption into a structured, predictable journey.

A structured approach typically involves these key steps:

- **Detecting technical debt** and assessing indicators of modernization readiness.
- **Inventorying data sources**, legacy ETL, integration workloads and governance tools.
- **Classifying use cases** by complexity, business criticality and migration readiness.
- **Identifying use cases apt for modernization** to leverage cloud-native capabilities — such as elastic scale, real-time processing and unified governance — and consolidating them into a unified cloud foundation.
- **Sequencing modernization in phases**, starting with high-value, lower-risk workloads to build momentum without overwhelming operations. These early wins help establish organizational confidence.
- **Retiring redundant tools** and prioritizing asset reuse to protect existing investments.
- **Maintaining governance and data consistency across hybrid environments** throughout the transition, rather than waiting until the migration is complete.

Taking Stock: Are You Ready to Modernize?

Signals That It's Time to Modernize

The decision to modernize rarely happens overnight. More often, it stems from a series of recurring frustrations that may feel manageable individually but collectively signal that your architecture has reached its limits.

If several of the following are true for your organization, the case for cloud modernization is clear and urgent:

- AI and analytics initiatives consistently stall at the data preparation stage
- GenAI and agentic AI pilots are unable to scale
- Data pipeline failures regularly require manual intervention
- Onboarding a new data source takes weeks rather than days
- Your team spends more time maintaining existing integrations than building new capabilities
- Data quality issues are discovered downstream by business users rather than caught at ingestion

- Governance and compliance policies are enforced inconsistently across platforms
- You cannot produce end-to-end data lineage without manual reconstruction
- Your current data infrastructure was not designed to handle real-time or streaming data at scale

No architecture scores perfectly against every measure. However, **if four or more of these signals resonate**, your technical debt has compounded to a point where incremental fixes can no longer close the gap.

Gartner predicts that “by 2027, 50% of business decisions will be augmented or automated by AI agents for decision intelligence”.¹¹ The urgency to become AI-ready is high. The question is no longer whether to modernize, but how to build the business case and choose the right path forward.

¹¹ Gartner Press Release, “Gartner Announces the Top Data & Analytics Predictions,” June 17, 2025.

From Reactive to Future-Ready

The enterprises pulling ahead in the AI era are not necessarily those with the largest data teams or the most sophisticated models. Instead, they are the ones that made a deliberate architectural decision to stop managing fragmented infrastructure reactively and start building a unified foundation proactively.

That decision is not without complexity. Migrating legacy workloads, consolidating fragmented tools and re-architecting for cloud-native scale is a significant undertaking. However, the window for competitive differentiation is narrowing. Every quarter spent maintaining brittle pipelines and working around governance gaps compounds technical debt and delays the AI readiness the business demands.

Leading enterprises are not waiting for the perfect moment. They are treating modernization as a strategic program — sequenced, phased and built on a platform that unifies integration, governance, quality and intelligence under a single architecture. By moving from point solutions to a unified, cloud-native and AI-ready data foundation, you can reduce cost and complexity while harnessing the power of AI faster.

As you look to the future, modernizing your on-premises workloads with Informatica's **Intelligent Data Management Cloud™** (IDMC) provides a modern, cloud-native data foundation. By unifying data integration, governance, quality, cataloging and **master data management** (MDM) at enterprise scale, IDMC delivers the AI-ready foundation that advanced analytics and agentic workloads require.

Informatica's **Cloud Modernization Program** helps you reach this goal faster and more safely by reducing migration cost, time, risk, complexity and effort through a unified, modern data infrastructure.

The path from reactive to future-ready is clearer than it seems. The question is: when will you start?

Next Steps

Wherever you are in your modernization journey, take your next step with these resources:



Just starting to build a case for modernization: Explore the **Cloud Modernization and Consolidation hub**



Running Informatica on-premises: Take the **PowerCenter Modernization Readiness Assessment** or **see what 300 PowerCenter users are saying about cloud modernization**



Evaluating MDM modernization: Take the **MDM Modernization Readiness Assessment**



Ready to see how others did it: Read how **Citizens Bank and other leaders achieved cloud modernization success**



Determine how much PowerCenter modernization could save you: Use the **PowerCenter Modernization Calculator**



Determine how much master data management modernization could save you: Use the **Master Data Management Modernization Calculator**



About Us

About Informatica

Informatica from Salesforce is a leader in AI-powered enterprise cloud data management. Its Intelligent Data Management Cloud (IDMC) platform enables organizations to connect, manage and unify AI-ready data across the enterprise. With capabilities spanning data cataloging, integration, governance, quality, privacy, metadata management and master data management, Informatica supports a broad partner ecosystem and helps customers unlock the full value of their data and AI initiatives.

About Salesforce

Salesforce is the #1 AI CRM, empowering companies to connect with their customers in a whole new way through the power of artificial intelligence, data, and trust. For more information about Salesforce (NYSE: CRM), visit: www.salesforce.com.

IN19-5345-0526

© Copyright Informatica LLC 2026. Informatica and the Informatica logo are trademarks or registered trademarks of Informatica LLC in the United States and other countries. A current list of Informatica trademarks is available on the web at <https://www.informatica.com/trademarks.html>. Other company and product names may be trade names or trademarks of their respective owners. The information in this documentation is subject to change without notice and provided "AS IS" without warranty of any kind, express or implied.

[informatica.com](https://www.informatica.com)

Where data & AI come to



Worldwide Headquarters
2100 Seaport Blvd.
Redwood City, CA 94063, USA
Phone: 650.385.5000
Fax: 650.385.5500
Toll-free in the US: 1.800.653.3871

[informatica.com](https://www.informatica.com)
[linkedin.com/company/informatica](https://www.linkedin.com/company/informatica)
[x.com/Informatica](https://www.x.com/Informatica)

CONTACT US