




Your data, *ready for the AI era*

How modern data orchestration turns AI ambition into enterprise-wide results.



The AI imperative has arrived.

Is your data ready?

Artificial intelligence is no longer a competitive differentiator, it is quickly becoming the baseline expectation for how modern enterprises operate. The question is no longer whether to invest in AI, but whether the data foundation is in place to make those investments count. The evidence is striking.

According to McKinsey's State of AI 2025 report, 88% of organizations now deploy AI in at least one business function, up from 78% just one year earlier. Yet despite this adoption, nearly two-thirds of organizations have not yet begun scaling AI across the enterprise. Most remain stuck in a cycle of pilots and proofs of concept that never fully transition into production.

An abstract background featuring a dense field of blue and purple particles, with some particles forming distinct, glowing trails or orbits, suggesting a complex data environment or a network of connections.

"Nearly two-thirds of respondents say their organizations have not yet begun scaling AI across the enterprise. *Only about 6% of organizations qualify as AI high performers.*"

McKinsey, The State of AI in 2025: Agents, Innovation, and Transformation

The bottleneck, in the vast majority of cases, is not the AI itself. It is the data. Models can only be as powerful as the information they can access, and in most organizations, that information is fragmented, siloed, slow to move, and difficult to trust.

Cirata exists to solve exactly this problem. Our data orchestration platform ensures that AI and machine learning initiatives are built on a foundation of data that is accessible, high-quality, and ready for action, wherever it lives, and wherever it needs to go.

Most AI projects *fail before they begin.*

The gap between AI ambition and AI outcomes is well documented. Projects get funded, teams get assembled, and models get built — yet the results rarely match the expectations.

The reason is almost always the same: the data infrastructure was not ready to support the demands of prod AI.

88%
of organizations use AI in at least one function

66%
have not yet begun scaling AI enterprise-wide

6%
qualify as true AI high performers with measurable EBIT impact


Source: McKinsey & Company, The State of AI in 2025: Agents, Innovation, and Transformation, November 2025.

The gap between early adoption and genuine enterprise impact is not a failure of ambition or talent. It is a structural problem: the data pipelines, migration tools, and orchestration capabilities that AI requires were simply never built for this purpose.

High-performing AI organizations — those achieving measurable EBIT impact from AI — share a consistent characteristic. They invest as heavily in data infrastructure as in model development. They treat data readiness not as a precondition to be solved eventually, but as the strategic foundation on which everything else depends.


The foundation every *AI initiative depends on.*

Making data AI-ready is not a one-time migration. It is an ongoing operational capability, the ability to ensure that the right data is always available, in the right place, at the right time, in a form that AI systems can use. This requires four core capabilities working together.



Real-time orchestration

AI models require live data, not yesterday's snapshot. Real-time orchestration automates and directs data flows across clouds and platforms as they happen — ensuring models always operate on current information.



Open & interoperable formats

Proprietary data formats create lock-in and slow down the movement of data between systems. Open formats like Apache Iceberg ensure that data remains portable, flexible, and usable across the full spectrum of ai and analytics platforms.



Seamless platform integration

Data infrastructure must connect effortlessly with the AI and analytics platforms your teams already use, without brittle custom integrations that require constant maintenance. Pre-built connectors to leading platforms eliminate the integration tax.



Enterprise-grade security & compliance

Moving large volumes of sensitive data introduces risk. Built-in security controls, full auditability, and compliance tooling ensure that data reaches ai systems through channels that meet regulatory requirements and organizational policies.

The platform that *turns data into AI fuel.*


Cirata Symphony is an intelligent data orchestration platform built for the scale and complexity that enterprise AI demands. Cirata Symphony moves and manages data across any environment, on-premises to cloud, cloud to cloud, or across regions securely, continuously, and without disruption to live systems.

Unlike legacy migration tools that require downtime, complex custom scripting, or repeated manual reconciliation, Symphony automates the entire process. It captures changes as they happen, ensuring that the data available to AI and analytics teams is always current and complete.


Freeing your AI teams to do what they do best

When data engineers spend their time maintaining fragile integration scripts, they are not building models.

Cirata Symphony eliminates the operational burden of data movement, redirecting your most valuable technical talent toward the work that actually drives business value.



What Cirata Symphony delivers.



Real-time orchestration

Continuous movement of data and metadata. Move datasets of any size from on-premises data lakes to any cloud environment without disrupting production workloads.



Zero disruption

Changes to source data are captured and replicated in real time without risk, code rewrites, or downtime, so models are never trained on stale information.



Freedom and flexibility

Native support for Apache Iceberg and other open table formats ensures data remains portable and future-proof.



Visibility and control

Full oversight of usage, cost, and security with automation that eliminates the need for hand-crafted scripts, freeing engineering teams to focus on building.



Enterprise trust

Single connective layer for mission-critical, regulated environments

What is standing between *your data and your AI goals?*

01. *Data silos and integration failures*

Enterprise data does not live in one place. It is distributed across legacy on-premises systems, multiple cloud platforms, departmental databases, and operational applications. This fragmentation is not just inconvenient, it is disqualifying for AI.

Training an effective model requires a unified, consistent view of data. Without it, models learn from incomplete information and produce unreliable outputs. The problem compounds over time: the longer organizations delay integration, the more entrenched the silos become.

02. *Slow and brittle data movement*

AI and machine learning models need large volumes of data, and they need it reliably. Slow data transfer between storage environments and analytics platforms creates bottlenecks that stall innovation cycles and frustrate data science teams.

Many organizations rely on hand-coded scripts or legacy replication tools that cannot keep pace with the volume and velocity of modern data, and that require constant manual intervention when something breaks.

03. *Poor data quality and consistency*

The adage 'garbage in, garbage out' has never been more consequential than in the context of AI. Inaccurate, incomplete, or inconsistent data does not just produce bad model outputs, it erodes organizational trust in AI as a whole.

When business leaders cannot rely on AI-generated insights, adoption stalls regardless of how powerful the underlying technology is. Ensuring data quality requires not just better tooling but continuous, automated validation at every stage of the pipeline.

Practical applications *across the enterprise.*

The value of AI-ready data is not theoretical. Organizations that solve the data readiness problem unlock a broad range of high-value use cases that would otherwise remain locked in pilot mode. Here are the areas where the impact is most immediate.

Accelerating model training

Training effective AI models requires access to large, high-quality datasets, and the ability to refresh those datasets as the underlying data changes.

Cirata Symphony ensures that training pipelines draw from the most recent available data, reducing the time from data collection to model deployment and improving accuracy.

Powering real-time analytics

Operational AI are models that inform live decisions in areas like fraud detection, pricing, and personalization, requires data that is current to the minute, not the day.

Cirata Symphony's real-time orchestration capabilities ensure that the data feeding these systems reflects the latest available state of the business.

Retrieval-augmented generation

RAG architecture allows large language models to query live organizational data rather than relying solely on pre-trained knowledge. This unlocks a new class of AI applications, but it only works when data is continuously available and properly indexed.

Cirata Symphony provides continuous data transfer that RAG systems depend on to respond at speed and with accuracy.

Supporting AI governance / compliance

As regulatory requirements around AI and data privacy become more stringent, the ability to demonstrate where data came from, how it moved, and who accessed it becomes a compliance requirement.

Cirata Symphony's full auditability and data transfer capabilities provide the paper trail that governance teams need.



Begin your AI enablement journey *with Cirata.*

AI transformation does not begin with a model. It begins with a decision to build the data foundation that makes models possible. Organizations that make this investment early gain a compounding advantage, every new AI initiative benefits from infrastructure that is already in place.

Cirata works with organizations at every stage of the AI readiness journey, from initial assessment through full-scale production deployment. Our approach is practical, not prescriptive, and we start where you are and build toward where you need to be.

The background of the lower half of the page is a composite image. It features a dark blue and purple color scheme with glowing lines and dots, resembling a data network or a stylized city skyline at night. Overlaid on this are snippets of code in various colors (green, yellow, blue) and some abstract geometric shapes. A white rounded rectangle is positioned in the lower-left quadrant, containing the contact information.

Contact us to learn more about Cirata Symphony and how we can help you *move, modernize, and activate your data for AI.*

cirata.com/data-orchestration/solutions/ai-enablement