

ORACLE

Enterprise Data Management: The Trusted Data Foundation for Enterprise AI





Table of contents

Introduction	3
Enterprise data management: The missing ingredient	4
Benefits of an EDM data foundation for enterprise AI	5
EDM as institutional memory for enterprise AI	8
From AI hype to AI value	9
How Oracle can help	10

Enterprises are racing to deploy AI models that promise automation, rich new insights, and revenue growth. But there's a disconnect between promise and reality. Recent research reveals the hard truth.

95%

AI pilots fail to reach production.¹

60%

AI initiatives will fail by 2026 because of poor data quality.²

81%

Companies say they still struggle with AI data quality.³

It's not that AI technology doesn't work. It's that most organizations try to run with it before they can walk. They jump into deploying models without a solid data foundation.

A July 2025 MIT NANDA report found that 95% of organizations reported no return on their generative AI investments, despite total spending of \$30 billion to \$40 billion. Gartner recently predicted that more than 40% of agentic AI projects will be canceled by the end of 2027 because of "escalating costs, unclear business value, or inadequate risk controls." Every acquisition, product launch, and reorganization changes a company's data landscape. Without a system to keep that data accurate and aligned, AI loses its effectiveness. Decisions fragment. Reports conflict.

Why do so many AI projects falter?

There are several root causes, including the following:

- **Gaps in data quality.** Incomplete, duplicate, or inconsistent data erodes model performance and trust.
- **Siloed and disorganized data.** Scattered information systems prevent AI from seeing context and connections.
- **Opaque data lineage.** Business leaders and regulators demand to know where the data came from, but most AI projects can't show that.
- **Outcome misalignment.** Too many AI projects start as a solution looking for a problem without goals for achieving measurable business results.

Enterprise data management: The missing ingredient

Enterprise data management (EDM) is the practice of organizing, aligning, and governing data at every stage of its journey across an organization, helping ensure its quality and consistency, even across disparate applications and departmental silos. Proper EDM is key to maximizing the business value of AI. In a finance context, the many benefits include improving financial forecasting; delivering faster, deeper analyses of financial data; identifying company growth opportunities; reducing the time it takes to close the books; producing reports more quickly; automating manual departmental tasks; and reducing the time it takes to audit financial statements.

AI innovation succeeds only when it's built on trusted data. Yet most enterprises struggle with data inconsistency, fragmented ownership, and governance fatigue, causing high AI failure rates and limited adoption.

Oracle Fusion Cloud Enterprise Data Management addresses these challenges with the following six-component framework:

- 1 Unified information model:**
Harmonizes data definitions across horizontal and vertical domains to create a shared single source of truth.
- 2 Governed change management:**
Supports disciplined changes to core business data by involving the appropriate stakeholders, applying agreed-upon rules and approvals, and maintaining a clear history that explains what changed, why it changed, and who was involved.
- 3 Cross-domain connectivity:**
Links operational (ERP, supply chain, HR) and analytical (enterprise performance management, analytics, database) platforms for contextual relevance through relationships, associations, hierarchies, and data maps.
- 4 AI-embedded productivity:**
Uses agents for application onboarding, data profiling, data matching, natural language querying, generating change requests, and more.
- 5 Quality and lineage intelligence:**
Monitors and scores data similarity, completeness, accuracy, and anomalies in real time.
- 6 Trusted data products:**
Publishes governed, explainable data assets ready for AI and analytics.

Benefits of an EDM data foundation for enterprise AI



Master Data Foundation

Multidomain foundation to build AI trust



Intelligent Data Organization

Relate or classify data based on shared meaning or lexicon



Relationship Intelligence

Generate intelligence from structure: hierarchies, data maps



Hallucination Reduction

Trusted enterprise data as ground truth for AI models



Data Quality Assurance

Scoring for data completeness and accuracy



Explainable Data Lineage

Complete data provenance for transparent, auditable decisions

Specifically, here's how Oracle Fusion Cloud EDM can help organizations avoid the biggest AI pitfalls.



Master data foundation

Oracle Fusion Cloud EDM builds a shared, governed single source of truth across all data domains: customer, supplier, product, finance, and HR. It helps ensure that names, structures, and codes are consistent and traceable from source to system. Whether your CRM system says “GHC Intl.,” your ERP system says “GHCI,” or your data warehouse lists “Global Health Corp.,” your EDM platform maps those names to one master record. That means AI models don't have to guess—they're trained on trusted, reconciled data that truly reflects your business



Smart data organization

Oracle Cloud EDM semantic discovery enriches data with context so models can learn from relationships, not just rows. By marrying lexical insights about similar words and word fragments with their underlying meaning, EDM can be used to connect the dots in supporting intelligent data search, matching, mapping, and classification. For example, AI agents can automatically identify that an acquired company's “operating expenses—R&D” corresponds to “innovation costs” at the parent company because their descriptions match closely in meaning.



Relationship intelligence

In Oracle Cloud EDM, relationships between entities—such as how products roll up to portfolios, customers to regions, and accounts to business units—aren't just metadata; they're context. By curating and governing these hierarchies, EDM effectively builds the enterprise's internal knowledge graph. Agentic AI workflows can use this context to move beyond isolated data points to reason through the “why” and “how” that connect them. For example, an AI agent can infer that a supply issue in one region impacts revenue forecasts in another because the relational fabric is explicitly defined. This structure drives insights across connected business processes



Hallucination reduction

AI models hallucinate when they draw incorrect or misleading conclusions from unverified or conflicting inputs. Validated data sets ground AI models, reducing spurious or untrustworthy outputs. This means that an AI assistant doesn't have to guess whether customer “ACME MFG Ltd” and “ACME Manufacturing LLC” are the same entity. Oracle Cloud EDM grounds AI in validated, reconciled master data. By rooting generative and analytical AI models in verified, approved, and version-controlled facts rather than inferred guesses, it dramatically reduces hallucinations.



Data quality assurance

Poor-quality data—missing fields, outdated codes, duplicate entries—poisons AI models from the start. A machine learning model predicting customer churn, for instance, will fail if half its inputs are empty (“Country: N/A” or “Customer Tier: Unknown”) or mislabeled (“Tier: Gold” versus “Class: G”). Scoring for completeness and accuracy helps ensure AI is fed fit-for-purpose data only. Oracle Cloud EDM scores data quality for every record as it's mastered through match and merge or through deduplication. Records are assessed for completeness, consistency, and accuracy, and they're combined for downstream use. A supplier record missing tax ID or currency code will fail validation, while another with inconsistent region mapping (“NA” versus “North America”) can be flagged for review. By filtering out bad or incomplete data early, Oracle EDM helps ensure that AI models are “fed” only fit-for-purpose, context-rich data, producing predictions and recommendations that can be trusted.



Explainable data lineage

AI outputs can't be trusted if we can't trace the data behind its decisions. Oracle Cloud EDM provides this foundation through a multilayered lineage model that captures how enterprise data is proposed, approved, structured, and consumed.

Node links are generated to trace the provenance of master data records to their original sources. When change requests are committed, they can trigger subscription requests in a downstream context. Request lineage reveals the business intent and explains who initiated a change, why it was made, and what approvals or policies governed it.

Viewpoint lineage shows which structures, hierarchies, and reporting perspectives were active at any point in time, helping ensure that AI systems can interpret data in the proper organizational context. Node lineage traces individual dimension values across their lifecycle—creates, updates, merges, reclassifications, and future-dated updates—so AI models can see how entities evolved rather than treating them as static.

Together, these lineage layers give AI the context essential to answering “why” questions: Why did revenue spike after a cost center realignment? Why did forecasts shift after a product hierarchy change? Why did a supplier risk score rise after it was acquired by another company? By giving AI full transparency into the provenance and evolution of enterprise data, Oracle Cloud EDM transforms opaque model outputs into explanations grounded in governed history, making decisions auditable, traceable, and worthy of trust.



EDM as institutional memory for enterprise AI

Enterprises often talk about “institutional memory”—the accumulated understanding of how the organization has evolved, why decisions were made, and what patterns shaped today’s operating reality. In AI systems, we use similar language: short-term memory for immediate context, long-term memory for enduring truths, and preference memory for how people or systems behave. Yet with every change event, most enterprise applications erase their own past and retain none of this institutional intelligence.

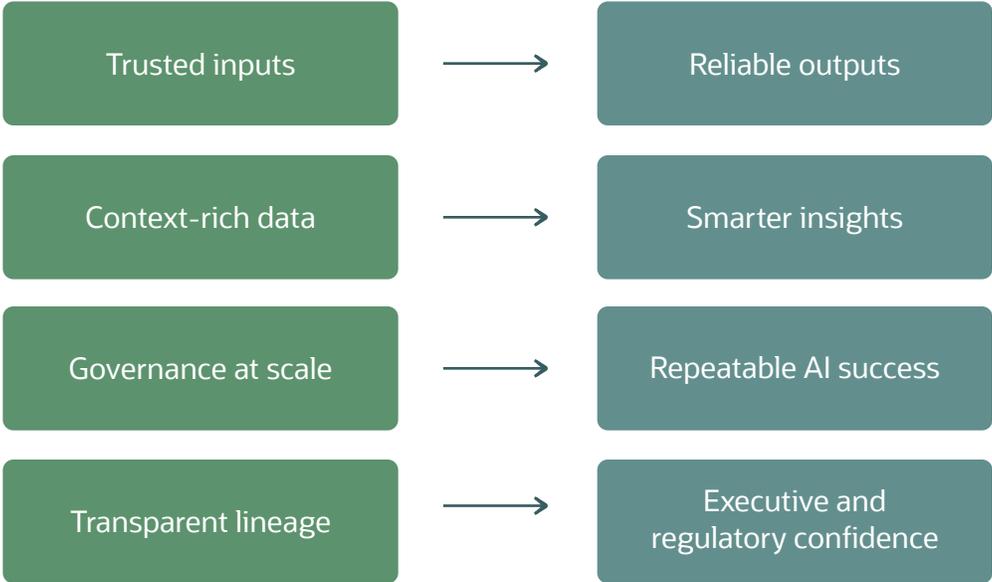
This is where EDM changes the equation. By versioning both data and metadata across time, EDM becomes the enterprise’s long-term memory: a durable, governed record of every structural shift, preference change, and contextual milestone. AI systems can finally move beyond analyzing static snapshots and instead learn from our “footprints in the sands of time”—the chronology of cost center migrations, customer preference updates, product launches, and future-dated plans that reveal intent, causality, and business trajectory.



Leveraging enterprise data past, present, and future for enterprise AI.

From AI hype to AI value

The message is clear: AI is not a shortcut around data problems. Without trusted data, organizations waste resources and lack confidence in their AI-informed decisions. With EDM, they create the conditions for AI to accelerate innovation and deliver measurable outcomes.



How Oracle can help

AI's future isn't just about bigger models; it's about better data foundations. Oracle Fusion Cloud EDM delivers exactly that: curated master data, validated relationships, and explainable lineage that transforms unreliable data into enterprise AI's grounding truth.

If 95% of AI pilots are failing, the way forward isn't another model. It's EDM. Don't wait, get your data ready for enterprise AI success today.

[Request a demo](#)

[Learn more](#)

Connect with us

Call +1.800.ORACLE1 or visit [oracle.com](https://www.oracle.com)

Outside North America, find your local office at [oracle.com/contact](https://www.oracle.com/contact)

1. Aditya Challapally, Chris Pease, Ramesh Raskar, and Pradyumna Chari, "The GenAI Divide: STATE OF AI IN BUSINESS 2025," MIT NANDA, July 2025.
2. "Lack of AI-Ready Data Puts AI Projects at Risk: Q&A with Roxane Edjlali," Gartner® Press Release, February 26, 2025. GARTNER is a trademark of Gartner, Inc. and/or its affiliates.
3. "Data Quality Is Not Being Prioritized on AI Projects, a Trend that 96% of U.S. Data Professionals Say Could Lead to Widespread Crises," Qlik, March 12, 2025.

Copyright © 2026, Oracle and/or its affiliates. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission. Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

