



# From Pilot to Production

The AI-Readiness Blueprint  
for the Hybrid Data Estate

# Contents

- 3** Introduction: AI Progress and the data modernization imperative
- 4** The five shifts toward an AI-ready data estate
- 6** A practical plan for chief data officers
- 7** Rocket Software: your partner in modernization
- 7** Building the foundation for future innovation



# Introduction

## AI Progress and the data modernization imperative

Enterprise AI is advancing at a remarkable pace, but progress often stalls when it meets the realities of production environments. The journey from promising AI pilots to trusted, production-scale applications hits a wall built from data access limitations, fragmented governance, and a fundamental lack of trust in the data itself.

For Chief Data Officers (CDOs), the pressure to overcome this challenge is immense. You're tasked with operationalizing AI and driving IT modernization—now a top-tier C-suite priority across industries—by providing trusted, governed data products at scale.

**The bottom line is clear:** AI's potential can only be reached with a solid data foundation.

Without it, even the most advanced models produce unreliable results, risk non-compliance, and fail to deliver business value. Leaders in data management understand this reality. They prioritize data quality, governance, and compliance, and they're actively deploying AI-ready infrastructures.

This focus enables them to outperform their peers in productivity, innovation, and revenue generation. To bridge the gap between AI ambition and reality, your data management team must fundamentally rethink its approach to data estates. AI doesn't just consume data issues—it amplifies them. A quality problem that used to affect a single dashboard gets repeated at scale, embedded into workflows, and delivered with confidence. This requires a strategic pivot toward a modern data architecture that's not only prepared for AI but designed to empower it.

This whitepaper outlines five critical shifts that will transform your data estate into one that's truly AI-ready. By embracing these principles, you can move beyond promising but ultimately frustrating pilots and start deploying trusted, scalable production AI solutions that drive real business outcomes.



# The five shifts toward an AI-ready data estate

Building an enterprise capable of maximizing AI requires more than new tools—it demands a cultural and architectural evolution. Here's a blueprint for creating a data estate that's robust, governed, and optimized for the age of AI.

## Shift one

### Establish trusted data with service-level agreements

For AI to be reliable, the data it consumes must be trustworthy. Trust, however, can no longer be a subjective debate or a retroactive quality check. It must become a measurable, non-negotiable release requirement, much like uptime or performance. An AI-ready data estate treats trust as a service-level agreement (SLA).

This means every production use case ships with defined Trust SLAs and the instrumentation to prove them. This is built on:

- **Named, accountable owners:** Who is responsible for critical datasets and AI outputs?
- **Accuracy:** Is the data correct and fit for its purpose?
- **Completeness:** Does the data provide full coverage without gaps?
- **Freshness:** Is the data timely and relevant for real-time decisions?
- **Pedigree:** Can you trace the data's lineage and sources with confidence?
- **One shared scorecard:** So everyone debates less and ships more.

Instead of relying on generic data quality dashboards, leveraging metadata makes trustworthiness, ownership, and lineage provable from the start. When CEOs ask for reliable AI copilots and predictive insights, it's the Trust SLA that makes deployment possible.

## Shift two

### Unify data and content into a single knowledge layer

AI models, particularly Large Language Models (LLMs), deliver their best results when they have access to the full context of your organization's knowledge. This includes both structured data (from databases and applications) and unstructured content (from documents, reports, and communications). Operating these as separate, siloed streams creates a fractured picture and limits AI's effectiveness. For instance, Retrieval-Augmented Generation (RAG) models fail when they can't access unified context.

Creating one governed knowledge layer that converges data and content is the second shift. This involves building a reusable, permissions-aware retrieval pattern that delivers cited, current, and consistent answers across all information sources. This approach solves the problem of "bot sprawl," where different bots give different answers. It reduces errors by ensuring AI responses are based on one complete and governed view of enterprise knowledge, ultimately enabling a single approach that can scale beyond one department.



## Shift three

### Implement sovereignty by design

Data sovereignty has traditionally been defined by where data is physically stored. In the age of AI, this definition no longer serves your needs. True sovereignty is about governing how data is used, not where it sits. AI introduces new complexities, requiring you to enforce purpose limitations, manage model access, and maintain provable audit trails.

In practice, sovereignty by design shifts the focus from storage location to evidence-backed enforcement. This is achieved by:

- Implementing a policy and audit plane that governs AI access at runtime
- Defining purpose limits, data masking protocols, and access approvals, ensuring data is used for approved purposes, not unlimited reuse
- Logging how every model and agent interacts with data, treating them as users with auditable identities and behaviors
- Making audit provability a feature, not an afterthought

This ensures that AI usage is a transparent process that aligns with regulatory requirements and internal policies to lower compliance friction and speed approvals for production deployments.

## Shift five

### Use machine learning for deterministic outcomes

While LLMs are powerful generators of content and ideas, their probabilistic nature can introduce risk and unpredictability in mission-critical operations. In short, they can't determine the accuracy or efficacy of their responses. To build stable and observable AI systems, you should use machine learning (ML) as a control system to provide guardrails.

In this model, ML provides deterministic outcomes—such as risk scoring, anomaly detection, routing, and verification—that guide and validate the outputs of LLMs. This creates a scalable, low-risk AI framework where responses and actions are predictable and implicitly trustworthy. By using ML as a control loop, your organization can shift its focus from simply building models to operating AI systems that are safe, monitored, and optimized for deterministic results.

## Shift four

### Build an agent-ready architecture

The future of enterprise automation lies with AI agents—autonomous systems that can perform complex tasks and workflows. However, simply deploying chatbots isn't enough. For agents to be reliable and effective, they need a specific architectural foundation. Without it, they risk using incorrect metrics, taking unsafe actions, delivering incorrect answers, and creating fragile automations.

An agent-ready architecture is built on four key components:

1. **Clear semantics:** Governed definitions and metrics that agents can understand
2. **Reusable context:** A clear view of entities and their relationships
3. **Governed tools:** A defined set of allowed actions and operations
4. **Domain-owned data products:** Stable, reliable interfaces for interacting with data

By developing this architecture, business units can publish governed agents for their specific data products. This moves interactions beyond simple query-and-response, enabling durable, safe, and sophisticated automation across your enterprise.



# A practical plan for chief data officers

Transforming your data estate is a significant undertaking, but you can make progress quickly with a focused approach. Don't try to boil the ocean. Here's a practical next-steps plan for CDOs to start executing on these five shifts quickly and move from AI pilots to trusted production.

## 01

### Define Trust SLAs for one high-value use case rather than 10 different ones

- Identify a critical AI initiative currently in the pilot stage
- Collaborate with business stakeholders to define specific, measurable Trust SLAs and clear owners for the data powering this use case. Focus on accuracy, completeness, freshness, and lineage
- Instrument the data pipeline to monitor and report on these SLAs, making trust a tangible metric

## 02

### Create a unified retrieval plan

- Map the structured data and unstructured content sources relevant to your chosen use case
- Plan a unified retrieval pattern that can access both data types through a single, governed interface
- Prioritize permissions-aware access to ensure the right users and models get the right information

## 03

### Implement sovereignty controls

- Establish a baseline policy and audit plane for AI data access
- Define and apply purpose limitations and data masking rules for your first use case
- Ensure all AI interactions with the data are logged to create an auditable trail, demonstrating sovereignty in practice

## 04

### Develop your first agent-ready data assets

- Based on your use case, define the core semantics (metrics) and context (relationships) for the key data product that drives it
- Design a stable interface for this data product and specify the governed tools (actions) an AI agent would be permitted to use

These initial steps will lay the foundation for a scalable, agent-ready architecture.



# Rocket Software: your partner in modernization

We understand that modernization is a journey, not a destination. With 35 years of experience helping thousands of customers, including 750 partners, run and modernize their mission-critical systems, we meet you wherever you are on your path. Our approach is guided by empathy and a human-first commitment to delivering with purpose.

The Rocket data portfolio is uniquely positioned to help you implement the five shifts for an AI-ready data estate.

- **Rocket® DataEdge™** enables you to prepare and govern data from structured sources, ensuring it's ready for analytics and AI
- **Rocket® ContentEdge™** provides the same level of control for your content repositories, creating a path toward a unified knowledge layer
- **Vertica™** offers a powerful analytics database to drive insights from your integrated data\*
- **Rocket® EVA™** is our agentic AI platform that delivers precise, end-to-end operational diagnostics across core systems

Together, these solutions empower you to move from operational systems to analytics outcomes, delivering trusted and unified data and content for your most critical risk, compliance, and AI initiatives.

[Learn more about our Data Solutions.](#)

## Building the foundation for future innovation

The promise of enterprise AI is immense, but its realization depends entirely on the strength of the data estate that drives it. The pilot-to-production wall is a real and persistent obstacle, but it's not insurmountable. By making the five strategic shifts—from establishing Trust SLAs to using ML for deterministic outcomes—you can break through this barrier. The key is to approach them in a measured way: define and instrument what "good" looks like, prove impact quickly, and then scale the pattern domain by domain.

These transformations build a foundation of trusted, governed, and unified data that not only reduces risks in AI adoption but accelerates it and amplifies its impact on your business. An AI-ready data estate enables you to deploy reliable copilots, create durable automations, and unlock new revenue streams with confidence. Your modernization journey today is building the foundation for AI-driven capabilities that will unlock even greater value tomorrow. With the right strategy and a committed partner, you can move your AI initiatives from pilot to production and secure a competitive advantage in the years to come.

[Talk to an expert](#)



\*Rocket has entered into a definitive agreement to acquire Vertica from OpenText.

## About Rocket Software

Rocket Software is a global technology leader in modernization and a partner of choice that empowers the world's leading businesses on their modernization journeys, spanning core systems to the cloud. Trusted by over 12,500 customers and 750 partners, and with more than 3,200 global employees, Rocket Software enables customers to maximize their data, applications, and infrastructure to deliver critical services that power our modern world.

Rocket Software is a privately held U.S. corporation headquartered in the Boston area with centers of excellence strategically located throughout North America, Europe, Asia and Australia. Rocket Software is a portfolio company of Bain Capital Private Equity.



Modernization. **Without Disruption.**™

Visit [RocketSoftware.com](https://RocketSoftware.com)

© Rocket Software, Inc. or its affiliates 1990–2025. All rights reserved. Rocket and the Rocket Software logos are registered trademarks of Rocket Software, Inc. Other product and service names might be trademarks of Rocket Software or its affiliates.

MAR-17716\_WP\_AI-Readiness-Blueprint\_V2

