SYNCARI

Governing AI Agents: A New Mandate for the Enterprise CDAO

By Nick Bonfiglio and Neelesh Shastry, Founders of Agentic MDM



As AI agents evolve from prototypes to autonomous actors across enterprise systems, CDAOs face a critical challenge: how to govern AI agents with the same rigor applied to data. Agentic MDM establishes the control plane for doing exactly that, ensuring agents act safely, accountably, and in alignment with enterprise policy.

1

Start with Governance, Not Just Data

AI governance must go beyond model risk and bias, into oversight of the agents themselves. Apply MDM principles like versioning, provenance, and policy enforcement to agent behavior. Define clear ownership, controls, and escalation paths now to avoid compounding risk.

2

Designate Ownership of Agent Actions

Autonomous doesn't mean unaccountable. Each agent needs a business owner empowered to:

- Oversee actions, approve changes, and respond to policy updates
- Ensure alignment with audit, compliance, and risk protocols

3

Embed Human Oversight Where It Matters

Not all decisions can be automated. Define review thresholds based on risk, data sensitivity, or regulation, and automate escalation for unresolved or high-impact issues.

4

Rethink Access Models for AI-Scale Execution

Centralized RBAC/ABAC governs agent access. Consider agent-specific policies that:

- Grants scoped, purpose-built visibility
- Routes traffic through a centralized orchestration layer (think: SSO for agents)

5

Make Agent Behavior Fully Auditable

Every agent action must be logged, versioned, and traceable. Ensure:

- Time-stamped change logs with context
- Version control for models, prompts, and policies
- Integration into audit and observability systems

6

Contain Governance Gaps at the Access Layer

Each uncontrolled agent connection (e.g., direct LLM-to-source) is a governance blind spot. Centralize access at the orchestration layer to reduce endpoint sprawl, enforce consistent policies, and simplify auditability.

The Role of Agentic MDM

Agentic MDM is more than infrastructure, it's the governance fabric that enables secure, scalable agent execution. It works alongside identity, observability, and lifecycle tooling to form a complete AI agent stack.

Final Thought:

AI agents are no longer a future state, they're operational. And no, a data warehouse isn't going to solve this for you - magically or otherwise. Your governance strategy will decide whether they become a source of advantage or an unmanaged risk.

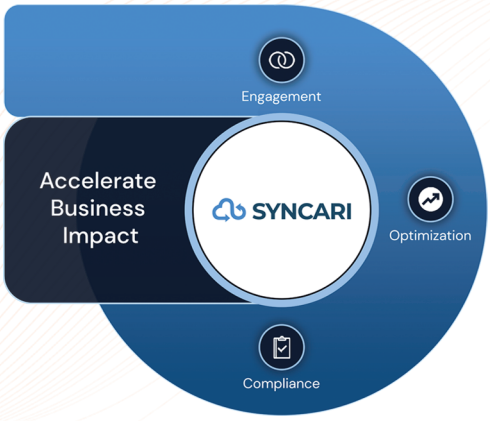
Syncari Agentic MDM™ is that control plane, to lead this transformation responsibly, securely, and at scale.

SYNCARI

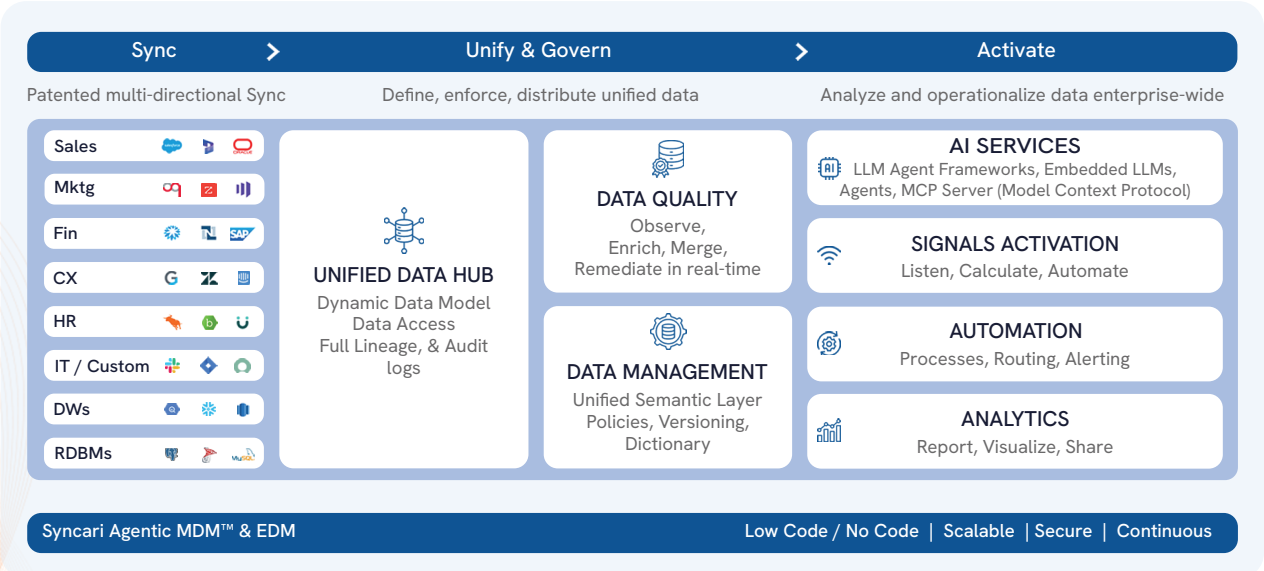
Data That Scales at the Pace of AI

One master data foundation with infinite multi-domain possibilities. Power operations, analytics, AI, and smart synchronization with Syncari Agentic MDM™.

- **Multi-Domain Mastering**
  Unify, standardize, and govern data across multiple domains—customers, vendors, products, finance, and more.
- **Dynamic Data Model**
  Adapt and evolve your data model with schema flexibility, custom entities, and hierarchical relationships that scale with your business.
- **AI-Ready Data Foundation**
  Establish a unified contextual data layer to power MDM and agent-driven workflows with enriched, real-time insights.
- **Intelligent Data Quality**
  Automate profiling, validation, remediation, and governance across multiple domains for AI-ready, trusted data.
- **Patented Multi-Directional Sync**
  Ensure real-time, consistent data across all connected systems without manual reconciliation.
- **APIs & Extensibility**
  API-first, fully programmable, with SDKs, HTTP, webhooks, and web services for seamless integrations.



Syncari Platform



Trusted by

Monotype. ALTAIR Antech Red Hat Trimble SERB Pharmaceuticals