

Agentic Master Data Management

Operationalizing Trust and Scale through
Autonomous AI Agents

Raghuvaran Kalluri
Lead Developer & Solution Architect
Royal Bank of Canada – US Wealth Management

The Breaking Point of Traditional MDM

- ⚠️ **Brittle, deterministic matching rules:** Matching fails when data schemas or semantic meanings shift.
- 📦 **Overwhelming stewardship queues:** Thousands of partial matches clog manual review processes.
- 🛡️ **Reactive governance:** Data quality checks happen post-ingestion, not proactively.
- 👤 **Inflexible hierarchies:** Inability to parse complex, unstructured corporate entity relationships automatically.



The Agentic Paradigm Shift

Traditional MDM (Deterministic)	Agentic MDM (Autonomous)
Relies on strict string similarity (e.g., Jaro-Winkler).	Understands semantic context and real-world meaning.
Dumps ambiguous records into manual stewardship queues.	Autonomously resolves complex exceptions and edge cases.
Struggles to identify relationships without explicit keys.	Traverses external data to map dynamic corporate hierarchies.

Empowering AI agents to ask: *"Do these records represent the same real-world entity?"*


Agentic Ecosystem Integration

 **Orchestration Layer:** Sits between data sources and the Hub.

 **Specialized Agents:** Semantic, Quality, Compliance, and Resolution act in concert.

 **HITL Escalation:** Ensures safety before updating the Golden Record.

Agentic MDM Architecture



The Agentic Operational Loop



Observe & Interpret

Agents continuously inspect current data states across diverse systems, understanding objectives within a broader enterprise context.



Decide & Act

Agents determine the next action (correct, enrich, escalate) and execute using authorized tools entirely within governance boundaries.




Explain & Continue

Every action is fully explainable, logged, and reversible. Agents persistently monitor outcomes and adapt their models.

Pillar 1: Transforming MDM with Autonomous Agents

-  **Augment, Don't Replace:** Native engines continue to handle high-speed, high-confidence deterministic merges (70% volume).
-  **The Cognitive Layer:** Autonomous agents intercept the remaining 30% exception queue.
-  **Semantic Resolution:** Agents read unstructured notes and cross-reference news to resolve complex B2B hierarchies.
-  **Human-in-the-Loop:** Reserved strictly for true anomalies below confidence thresholds.



Pillar 2: Improving Data Quality & Governance



Shift-Left Governance

Moving quality checks from post-ingestion to the point of entry. Agents profile and cleanse data as it streams in.



Active Compliance

Agents continuously cross-reference clients against sanctions lists and adverse media before transactions occur.



Self-Healing Pipelines

Agents detect schema drift or completeness drops, autonomously adjusting mappings or halting workflows to prevent contamination.


Pillar 3: Enabling Trusted, Scalable AI

-  **The Prerequisite for GenAI:** Agentic workflows fail rapidly when acting on hallucinated or fragmented enterprise data.
-  **Context-Aware Synthesis:** Moving beyond a monolithic Golden Record to provide purpose-built views (e.g., Risk vs. Marketing lenses).
-  **Glass-Box Auditability:** Agents produce version-controlled logs detailing the exact reasoning for every merge to satisfy regulatory scrutiny.




The Knowledge Graph Foundation

-  **Relational Awareness:** Agents reason over relationships, lineage, and complex B2B hierarchy, not just flat tables.
-  **Policy Integration:** Governance rules and risk tolerances are hardcoded directly into the graph architecture.
-  **Historical Evidence:** Agents leverage past survivorship decisions and stewardship actions to constantly improve confidence scoring.



Adopting Agentic MDM



Questions?

Stop managing rules.
Start governing agents.