



SOLUTION BRIEF

## Data readiness AI agent

Make raw data AI-ready via mapping, cleansing, and entity reconciliation.

# 01

---

## Solution overview

**Service name:** Data readiness AI agent

**One-line outcome:** Make raw data AI-ready via mapping, cleansing, and entity reconciliation.

**Audience:** Chief Data Officers, Heads of Analytics, IT/Data Platform leaders, Data Engineering/ML Platform teams, and Business Operations owners (finance, risk, customer, supply chain).

**Partners:** Works with your existing data prep/wrangling stack plus major cloud platforms and leading data-quality, catalog, and observability vendors (e.g., dbt Labs, Monte Carlo, Telmai).

# 02

---

## Solution snapshot

### **What it does:**

The Data Readiness AI Agent maps raw data into a canonical schema, assesses quality, corrects issues, cleans/standardizes fields, imputes missing values where appropriate, and performs master data management (MDM) tasks like entity resolution and record reconciliation. It incorporates Gruve's Data Quality AI Agent for scoring, fix recommendations, approval workflows, and verified remediation as part of readiness execution.

### **Why it matters now:**

As organizations scale AI and advanced analytics, gaps in completeness, accuracy, and lineage create model risk, compliance exposure, and operational rework—making “AI-ready data foundations” urgent. Gruve positions its Data & AI services around enabling enterprises through each stage of the AI journey, which starts with trustworthy inputs.

**Time to value:**

- Go live in 4–8 weeks (pilot domain + canonical mapping + readiness pipeline)
- First value in 15–30 days (first canonical dataset + quality lift + reconciled entities)

**15-40%**

reductions in data prep effort/rework

**Typical ROI:**

- Teams often unlock measurable productivity by reducing data preparation and rework; industry studies cited in the Data Quality Review reference 15–40% reductions in data prep effort/rework and 10–20% faster reporting cycles.
- Data practitioners commonly spend large portions of time on cleaning/organizing data (e.g., ~60% reported in a widely cited survey), making automation and repeatable readiness pipelines a high-leverage investment.
- Poor data quality costs organizations ~\$12.9M/year on average, so preventing defects from reaching analytics/AI can produce multi-x payback.

**10-20%**

faster reporting cycles

**60%**

of data practitioners commonly spend large portions of time on cleaning and organizing data

# 03

---

## Core value

**Primary outcome:**

Deliver “AI-ready” datasets by converting raw sources into a canonical, validated, and reconciled representation—typically targeting measurable improvements such as higher completeness/validity, lower duplicate rates, and fewer downstream incidents within the first 30–60 days.

### **Secondary outcomes:**

- Efficiency: Less manual cleansing/reconciliation and faster delivery of analytics/AI features.
- Risk: Reduced reporting and audit risk through better-controlled pipelines and evidence-based quality controls.
- Experience: Higher stakeholder trust via transparent scores, issue catalogs, and lineage/verification artifacts.

### **Why Gruve:**

Gruve combines deep data engineering and AI implementation expertise with enterprise data strategy and governance experience—so readiness is delivered as an operational capability aligned to downstream AI/analytics and regulatory needs, not a one-off cleanup.

# 04

---

## Key benefits



**Faster AI/analytics delivery:** Less time lost to wrangling and rework; readiness pipelines produce consistent, canonical datasets for dashboards, models, and agents.



**Lower model and compliance risk:** Quality controls + traceable remediation reduce exposure from incomplete/inaccurate inputs and support evidence-based governance.



**Higher data trust and adoption:** Published readiness/quality scoring, issue logs, and verification artifacts increase confidence in outputs and speed stakeholder alignment.



**Durable improvements (not regression):** Readiness isn't "done once": the agent operationalizes tests/monitors and repeatable remediation patterns for sustained quality.

# 05

---

## How it works

### Tier 1/Feature 1: Canonical mapping + readiness baseline (pilot)

**Summary:** For teams needing quick traction, we select 1–3 high-value domains (e.g., customer, finance, risk), define a canonical schema, map raw sources into it, and establish a readiness baseline with prioritized remediation actions.

#### **Core features included:**

- Canonical schema definition (business + technical alignment)
- Source-to-canonical mapping (field/semantic mapping, validation rules)
- Readiness scoring: completeness/validity/consistency + key entity integrity checks
- Data Quality AI Agent integration: score → recommend fixes → approval workflow → execute fixes

#### **Typical use cases:**

- AI pilot kickoff where input reliability is a gating risk
- ERP/CRM/warehouse migration stabilization
- Post-merger source consolidation and definition alignment

**Key outcome or benefit:** A first “AI-ready” canonical dataset with measurable quality lift and an approved remediation backlog.

## Tier 2/Feature 2: Automated readiness at scale (MDM + continuous controls)

**Summary:** For organizations scaling across domains, we operationalize the agent to run continuously: ongoing quality remediation, standardization, missing-value handling (where appropriate), and MDM/entity resolution to maintain a single, reconciled view over time.

### Core features included:

- MDM/entity resolution (match/merge rules, survivorship, golden record patterns)
- Continuous DQ: dbt-style tests + observability monitors + anomaly detection where applicable
- Governance artifacts: change logs, stewardship workflows, KPIs and monitoring dashboards
- Expansion playbook: onboard additional domains/sources with repeatable accelerators

### Typical use cases:

- Enterprise AI programs requiring sustained input reliability and traceability
- Regulated industries where data quality compliance is mandatory
- Scaling self-service analytics/ML across many teams and products

**Key outcome or benefit:** A durable, scalable “data readiness” operating capability that continuously feeds AI and analytics with trustworthy, canonical, reconciled data.

## About Gruve

Gruve partners with leading enterprises to transform data into measurable business impact. Our team brings deep expertise in enterprise data architecture, AI and analytics strategy, cloud modernization, and organizational change. We combine technical rigor with business acumen, ensuring recommendations are both architecturally sound and executable within your organizational constraints. With proven success across financial services, healthcare, manufacturing, and technology sectors, Gruve delivers data and AI solutions that drive growth, efficiency, and competitive advantage.

Learn more at [www.gruve.ai](http://www.gruve.ai).