WHITEPAPER

Tiger Analytics™

BRIDGING THE AI READINESS GAP

How Enterprises Can Build an AI-Ready Data Foundation

For data leaders and engineers building agent-ready platforms

Table of Contents

Executive Summary

Every enterprise leader we speak with today is asking the same question:
How do we unlock measurable business value from AI?

Behind the boardroom optimism sits a harder engineering reality. AI systems perform well in demonstrations but often stall in production. Agents return inconsistent answers, retrieve outdated policies, and rely on stale data. The natural instinct is to blame the model. However, in most cases, the root cause is elsewhere. The root cause is often not the model, but the data foundation beneath it.

At **Tiger Analytics**, while partnering with various Fortune 500 companies across banking, retail, healthcare, and industrial manufacturing, we have seen the same pattern repeat. The primary obstacle to enterprise AI is the underlying data foundation, not the AI model itself. Platforms optimized for two decades of human analyst consumption cannot serve autonomous AI agents that reason, decide, and act without a human in the loop.

This paper serves three purposes. It defines the **AI Readiness Gap** and introduces a shared vocabulary for describing what is broken. It then presents the **Data Health Check** and the **Six-Dimension AI Readiness Model** as a diagnostic framework to help you assess your organization's current state. Finally, it outlines an architecture and capability roadmap tailored to organizations optimizing existing platforms or building new ones.

> ***The question is not how to apply AI to your data, but whether your data is AI-ready.***

What you will take away

- A shared vocabulary to identify and describe the AI Readiness Gap within your organization.
- A Data Health Check and Six-Dimension diagnostic to score where you stand.
- Two clear paths forward: Optimize (for legacy enterprises) or Leapfrog (for greenfield teams).
- A reference architecture, the Intelligence Data Lake, plus the Knowledge Layer in depth.
- A ten-capability roadmap sequenced across three phases.
- A view of how Tiger Analytics can partner with you on this journey.

1. Understanding the AI Readiness Gap

Agentic AI is no longer speculative. Production systems today are investigating sales drops, reconciling quarter-ends, and synthesizing compliance policies. In the right enterprises, they are collapsing hours of analyst work into minutes, with better audit trails than those produced by humans. The economics are compelling.

Yet for every deployment that scales, we see several stall in pilot. In our conversations with Chief Data Officers, Analytics Heads, and Engineering Leaders over the last eighteen months, the frustration follows a recognizable pattern.

What we hear from clients

> ***"Our agents give different answers depending on who asks.***
>
> *A financial services CDO on inconsistent metric definitions across three business units*

> ***"We spent six months on prompt engineering. It was the wrong problem***
>
> *A retail analytics leader whose agents were retrieving stale product catalogs.*

> ***"Compliance won't clear our agent for production. We can't explain how it reaches its answers.***
>
> *A bank building an internal reasoning assistant.*

> ***"We have a lake, a warehouse, a catalog. Why is this still hard?***
>
> *An engineering leader whose platform was mature by BI standards but not by agent standards.*

The pattern underneath

Four failure modes recur across almost every stalled deployment we have investigated. None of them are model failures.

- **Inconsistent metrics:** Three agents asked, "what is our churn?" return three numbers, because no semantic layer defines it once.
- **Identity confusion:** Cust_001 in CRM, customer_id_1 in ERP, and CUST-MUMBAI-01 in CDP are treated as three different customers.
- **Confident hallucination:** Agents act on stale data because freshness signals were never exposed at the consumption layer.
- **Shallow reasoning:** Agents answer single-table questions but fail at multi-hop questions, because the platform has tables but no graph.

Each of these is a data foundation problem dressed up as an AI problem. This is empowering, because it means the fix does not depend on the next foundation model release. It depends on choices your data engineering organization can make this quarter.

Why Traditional Data Platforms Are Not Designed for AI Agents

For the past two decades, data engineering has been optimized for a single consumer, the human analyst working with dashboards. That design philosophy shaped enterprise data platforms for years but is no longer sufficient for AI agents.

What was built for humans	What agents actually need
Dashboards for visual consumption	APIs for programmatic reasoning
Cryptic column names like rev_adj_net_q	Self-describing semantic schemas
Documentation in wiki pages	Machine-retrievable, versioned knowledge
Batch refreshes overnight	Event-driven, near-real-time access
Data quality signals in engineering dashboards	Trust signals inline on every result
RBAC for Analyst and Manager roles	Identity for autonomous, non-human actors

How can we define the AI Readiness Gap?

Fundamentally, it is a **structural problem** rather than a model or tooling issue. The AI Readiness Gap stems from the fact that existing data platforms were meticulously engineered for human consumption, while autonomous agents operate under an entirely different set of requirements.

At its core, the gap represents an opportunity for enterprises to evolve their data foundations to support agentic needs, specifically in two critical dimensions.

1. **The Semantic Layer:** Enterprises need to provide the domain context, such as domain-specific business rules, entity definitions, and the policy logic necessary for agents to perform accurate multi-hop reasoning.
2. **The Trust (Data Observability) Layer:** Data platforms need to provide real-time, inline trust signals (e.g., data quality scores, freshness metrics, lineage), which agents need to verify data accuracy and act autonomously.

For organizations looking to integrate autonomous agents into their enterprise, the challenge hinges on a pivotal structural evolution: moving beyond human-centric dashboards to engineer an Intelligent Data Platform rooted in knowledge and trust, specifically designed for agentic reasoning.

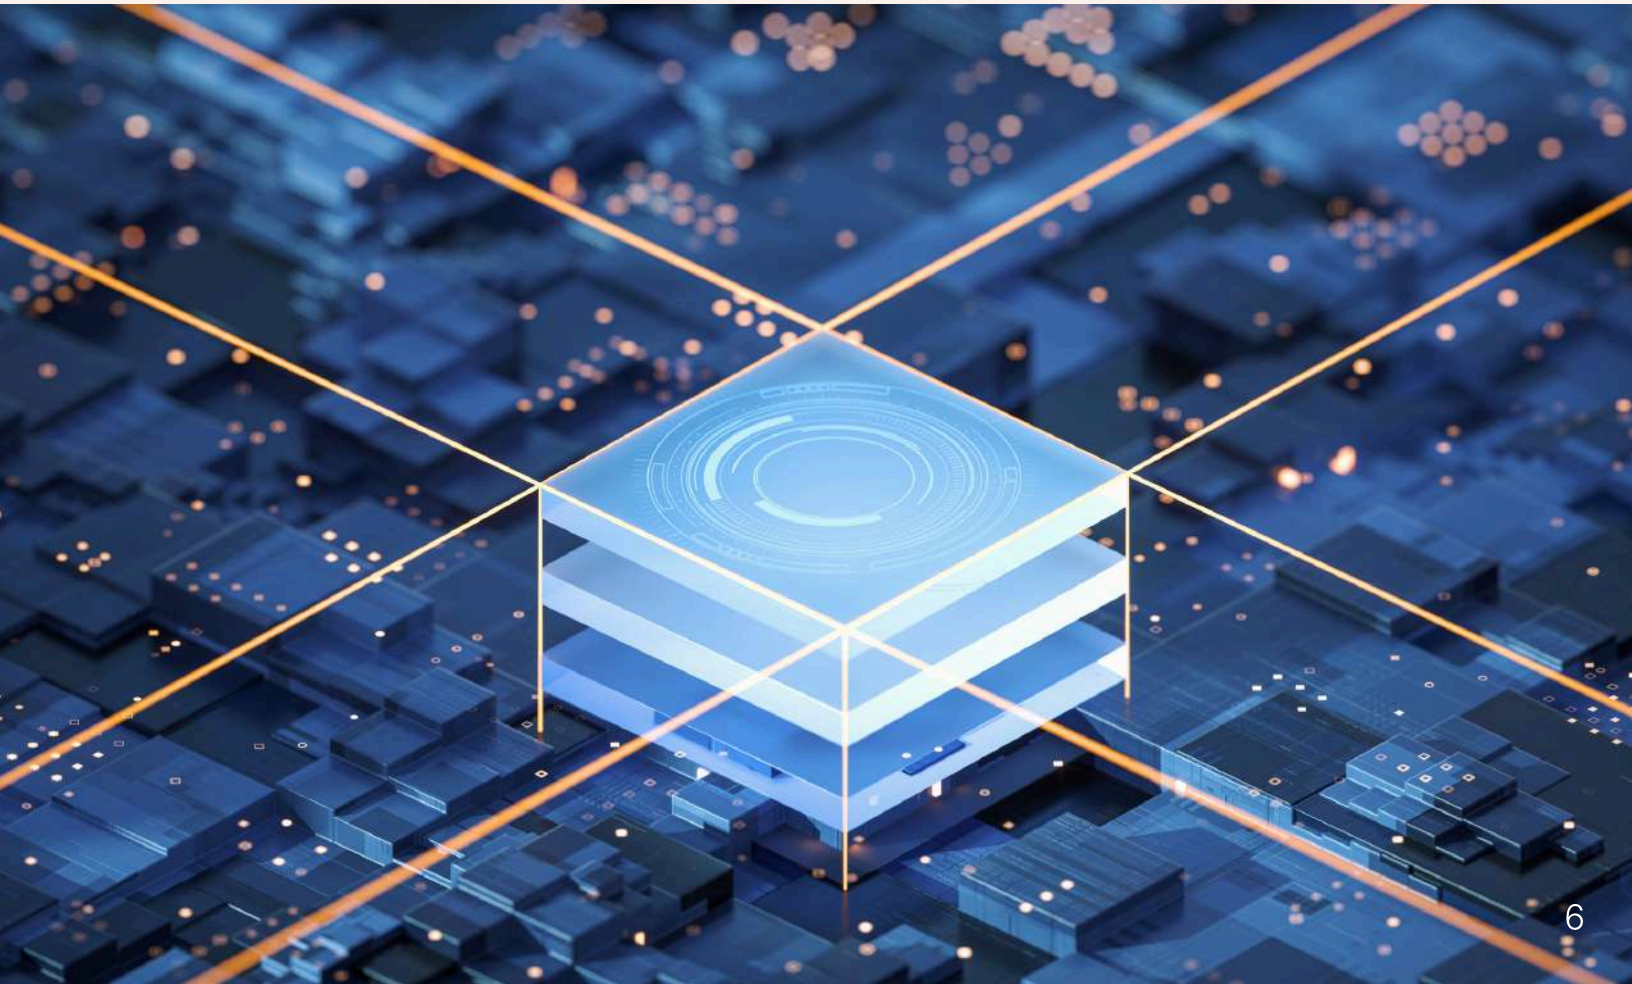

2. Assess: Where You Stand:The Diagnostic

Ask any leader in your organization, **"Is our data healthy?"** and the answer is likely to reflect BI-era assumptions, such as dashboard timeliness, pipeline uptime, and report freshness.

When we’re building for AI agents, data health takes on new definition. AI readiness is no longer defined by human-oriented data requirements. Instead, it requires enterprise data to be curated for agentic reasoning.

2.1 The five vital signs of data health for AI Readiness

Vital Sign	The question it answers	Traditional Check (Optimized for BI)	AI Readiness Check (Optimized for Agents)
Meaning	Do humans and machines interpret data the same way?	Business glossary exists	Machine-readable semantic layer
Identity	Do we know who a customer is across systems?	MDM project on the roadmap	Canonical entities in production
Freshness	How current does the consumer need the data to be?	Nightly batch is sufficient	Latency SLA per data product
Trust	Do we know what to believe?	DQ dashboard in a separate tool	Confidence and freshness provided inline
Coverage	Can we answer the questions the agent will ask?	Reports satisfy known needs	Ontology maps the unknown

How it works

A simple exercise for your leadership team:

Step 1: Pick your three most important data domains based on enterprise use case prioritization for agents (Customer, Product, Transaction, loyalty etc.)

Step 2: Score each dimension across these vital signs on a scale of 0 to -5.
In our experience, most enterprises score themselves 3 or higher on Freshness and Coverage, but only 1 or 2 on Meaning, Identity, and Trust.

That imbalance is where AI investments quietly lose value.

Understanding these imbalances is only the beginning. The next step is to assess where your organization stands today and identify the capabilities needed to close the gap.

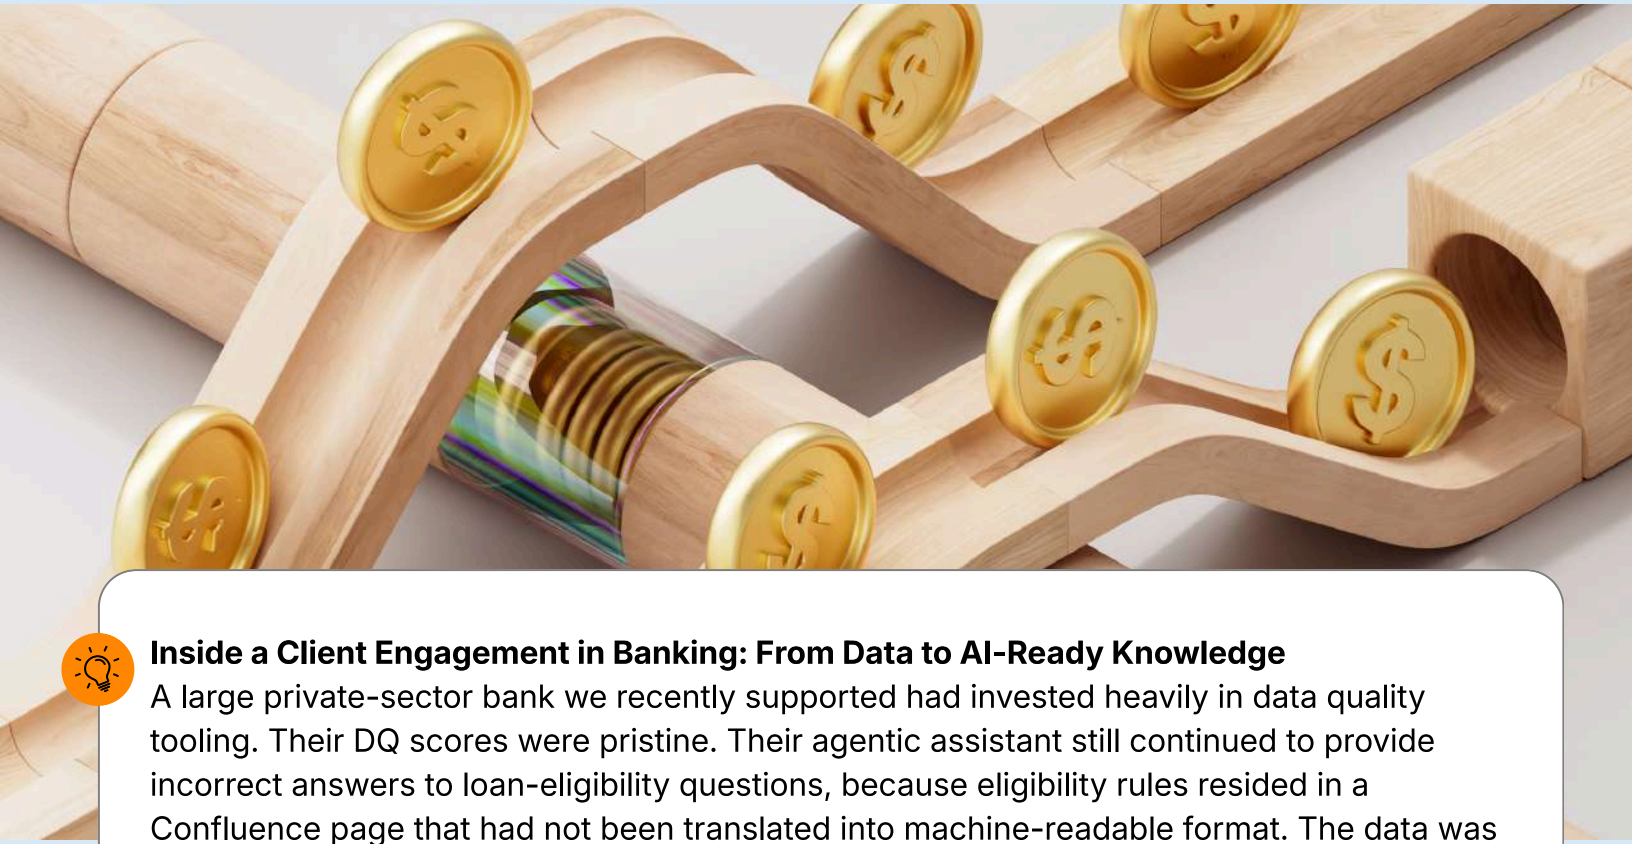

Inside a Client Engagement in Banking: From Data to AI-Ready Knowledge
A large private-sector bank we recently supported had invested heavily in data quality tooling. Their DQ scores were pristine. Their agentic assistant still continued to provide incorrect answers to loan-eligibility questions, because eligibility rules resided in a Confluence page that had not been translated into machine-readable format. The data was clean. Their knowledge was not.

If data health for AI is a different question, you need a different way to measure it. Two assessment tools help you determine where your organization stands. The Maturity Model indicates your current level of AI readiness, while the Six-Dimension Scorecard identifies the specific capability gaps you need to address.

Data Health check up is for AI readiness, and Data Platform maturity is for DP readiness

2.2 The Data Platform Maturity health check-up

Level	Stage	Reality	What we see in the field
L0	Data Chaos	Siloed, ungoverned, inconsistent	Rare in enterprises we work with; common in early-stage startups
L1	Data Available	Centralized and queryable, for humans	Where a large share of mid-market enterprises sits today
L2	Data Governed	Quality, lineage, catalog in place	Where most large enterprises with a mature BI program sit
L3	Data Semantic	Meaning encoded: metrics, ontology, glossary	Emerging: a small share of enterprises have crossed here
L4	Agent Ready	Machines can discover, trust, and reason	The target for agentic AI programs
L5	Self-Optimizing	AI improves the platform itself	Aspirational: the frontier

Inside a Client Engagement in CPG: L1 → L3

A CPG client operating across 12 South Asian markets sat at a solid L1. Data was centralized in a warehouse and queryable by their BI team. When they tried to deploy an agent for demand-sensing, it failed inside a week: three markets had different definitions of "active retailer," the SKU hierarchy differed by country, and no ontology captured the substitutability of similar products. A bigger model would not have helped; the fix was to build the Knowledge Layer for one region first.

Experiences like this are common across industries. The next step is to assess whether similar gaps exist within your own organization.

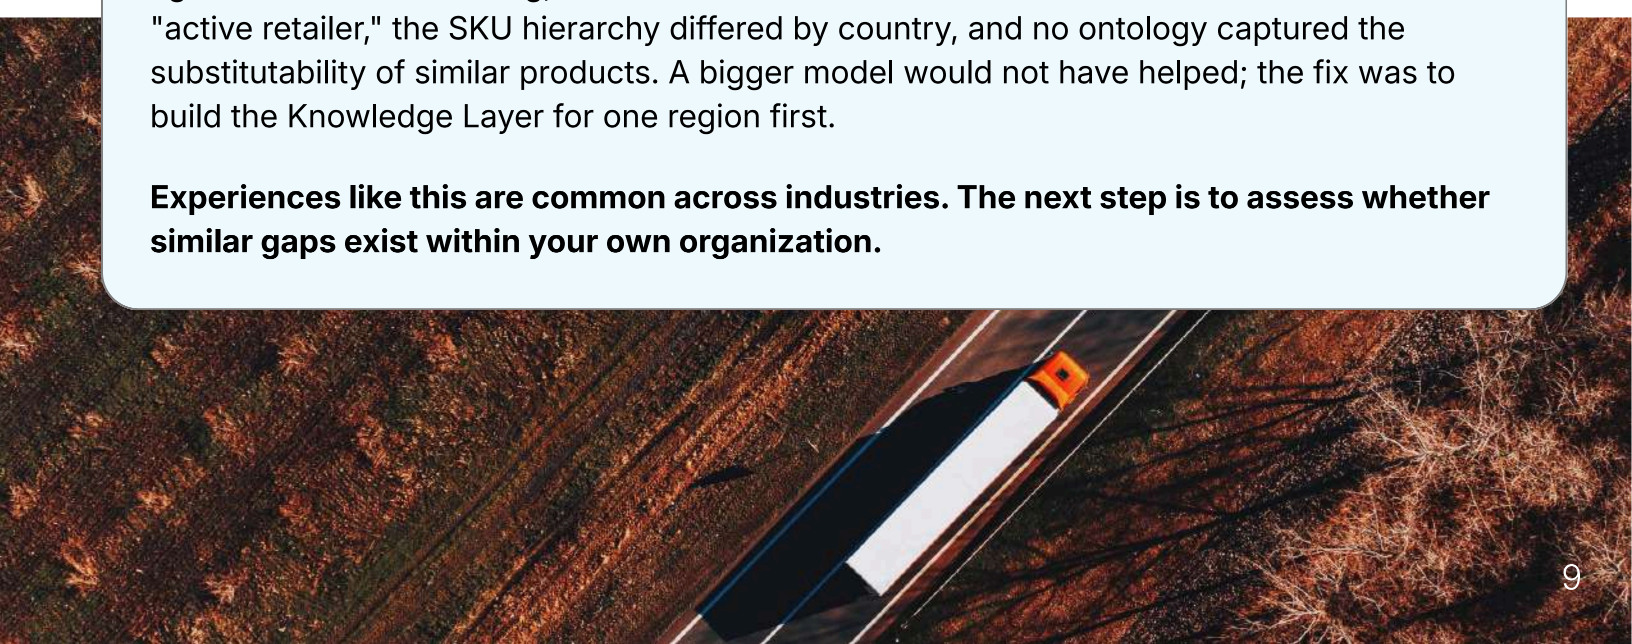

2.3 Agent Readiness Scoring model

After AI readiness Health check, you will understand if your data follows the necessary standards for agentic AI deployment. Data Platform Maturity model will help you understand how mature your data platform is and whether it provides the capabilities needed to support agentic AI deployment.

After completing the AI Readiness Health Check and Data Platform Maturity assessment, use the Six-Dimension Scorecard to identify the capabilities that need to be strengthened for agent deployment.

Score each dimension on a scale of 0 - 5 for your platform. The pattern most enterprises exhibit, with stronger scores on Dimensions 1 and 5 and weaker scores on Dimensions 2 and 4, makes the AI Readiness Gap measurable.

Dimension	Core question	Common state today	Mark the Score
1. Data Foundation	Can you find it, move it, version it?	Generally healthy	
2. Semantic Layer	Does data mean something to a machine?	Almost universally missing	
3. Access Layer	Can AI reach data safely and quickly?	Built for SQL, not agents	
4. Trust Layer	Can AI know what to believe?	Almost universally missing	
5. Governance & Security	Can you control what AI does?	Built for humans, not agents	
6. Operations	Can you sustain AI-driven systems?	DataOps mature, AgentOps nascent	

Radar view: The Shape Of The Gap

The example below illustrates a typical readiness profile, with stronger scores on Foundation and Governance and weaker scores on Semantic and Trust. While Foundation and Governance typically score well, Semantic and Trust layers remain consistently weak, with the remaining dimensions falling in between. The relative pattern, rather than the individual scores, is what matters most.

Dimension	Score (0–5)	Visual
1. Data Foundation	4/5	
2. Semantic Layer	1/5	
3. Access Layer	3/5	
4. Trust Layer	1/5	
5. Governance & Security	4/5	
6. Operations	3/5	

> ***Most enterprises score well on Foundation and Governance.***
> ***They almost universally fail on Semantics and Trust. That is the gap.***

3. Action: Two Paths Forward: Optimize Or Leapfrog

Your assessment results show where your enterprise stands and where the biggest capability gaps remain. The next step is deciding how to address those gaps.

Not every enterprise starts from the same place, and the path to AI readiness should reflect that reality. Across our client engagements, two distinct archetypes emerge. Recognizing which one best describes your organization is the first step toward selecting the right approach.

3.1 Path A: Optimize

- Build on existing investments rather than replacing them. Layer new capabilities onto the existing platform.
- Add the Knowledge Layer as a new tier above the warehouse: semantic definitions, ontology, vectors, and business rules.
- Retrofit trust signals into existing pipelines rather than rebuilding them.
- Introduce data contracts **across the ten most critical** producer–consumer pipelines first.
- Expected timeline: **12–15 months** to production-grade agentic AI in one domain.

Inside a Client Engagement: Optimizing Path in Practice: For a global insurance client with a mature Snowflake platform, we did not touch the warehouse. We layered a semantic model, built an ontology for policy and claims, and exposed structured tool-use APIs. Nine months in, their claims agent went into production with a full audit trail. The existing platform remained intact while new capabilities were layered on to support agent-ready workloads.

3.2 Path B: Leapfrog

For digital natives, greenfield business units, and enterprises spinning up new data platforms from scratch. Organizations building new platforms have the opportunity to avoid the constraints of legacy BI architectures and design for AI readiness from the outset.

- Design the Knowledge Layer alongside the data foundation instead of treating it as a later addition.
- Bake data contracts into your platform as a first-class primitive from day one.
- Choose storage and compute that natively support graph and vector, not just tabular.
- Build agent identity and observability into your IAM before your first agent goes live.
- Expected timeline: 6–9 months to production-grade agentic AI in one domain.

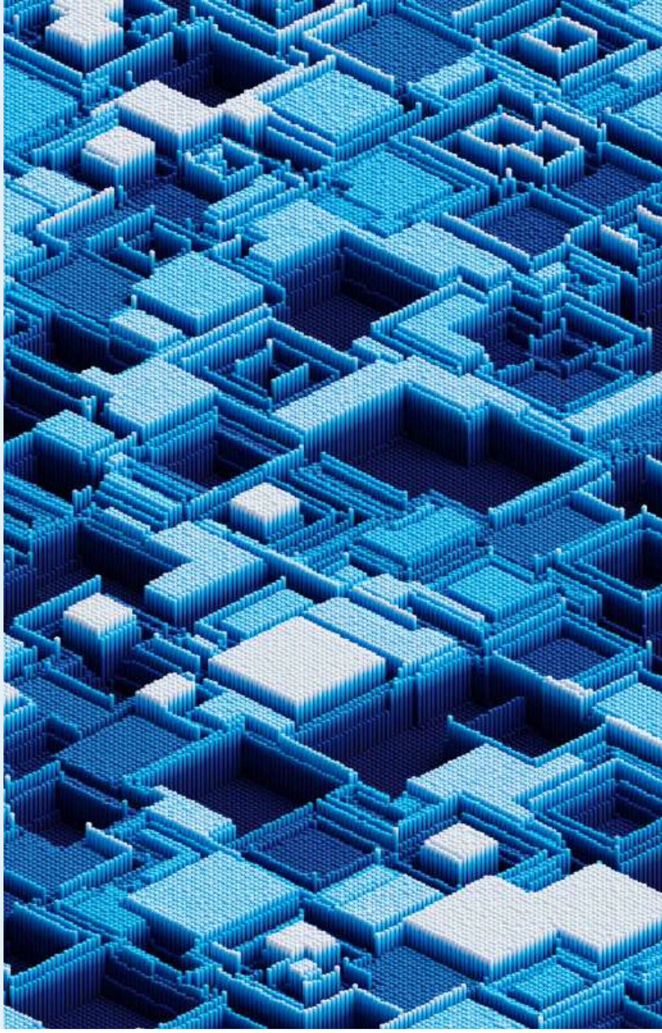

3.3 AI Readiness Deployment Matrix: Getting Your Data Agent-Ready

The following matrix summarizes the recommended agent deployment starting point based on an organization's current level of data maturity.

Your situation	Data Maturity	Recommended Path	Priority focus
Large legacy enterprise	L2 (Governed)	Optimize	Semantic Layer, Knowledge Layer overlay
Digital-native or startup	L1 (Available)	Leapfrog	Foundation + Knowledge Layer together
Regulated (BFSI, healthcare)	L2 with rich lineage	Optimize	Trust signals, governance, audit
Multi-country CPG or retail	L1–L2 across markets	Optimize	Entity resolution, domain ontology
Greenfield BU inside enterprise	L0–L1	Leapfrog	Contracts + semantic layer first
Mature analytics org	L2–L3	Optimize	Tool APIs, agent identity

Although organizations begin from different levels of maturity, both paths ultimately lead to the same outcome: an AI-ready platform.

4. Architect: The Intelligence Data Lake

An agent-ready platform consists of five architectural layers that connect enterprise data sources to AI agents. The first three layers represent the capabilities that many organizations have developed already.

The Knowledge Layer and Agent Interface extend the existing platform to support agentic AI deployment. Together with the first three layers, they form the Intelligence Data Lake.

ARCHITECTURE

The Intelligence Data Lake:

Conceptual Architecture

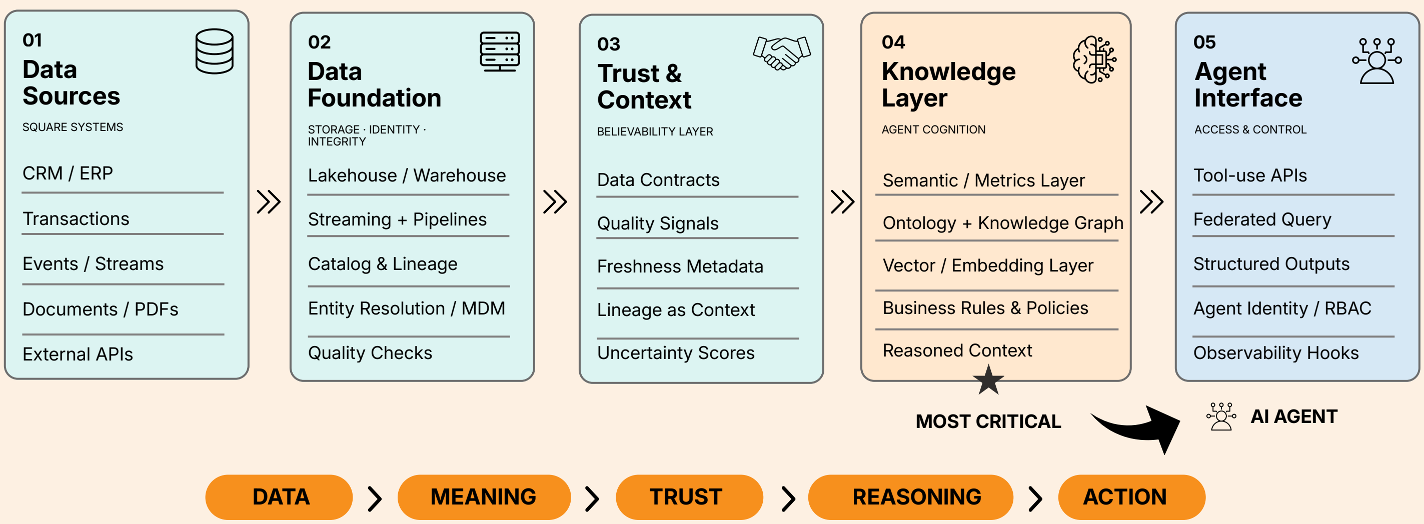

Stage	Owns	Key capabilities
1. Data Sources	Origin systems	CRM, ERP, transactions, events, documents, APIs
2. Data Foundation	Identity & integrity	Lakehouse, pipelines, catalog, lineage, MDM, quality
3. Trust & Context	Believability	Data contracts, quality signals, freshness, lineage as context
4. Knowledge Layer	Meaning & reasoning	Semantic layer, ontology + graph, vectors, business rules
5. Agent Interface	Safe access	Tool APIs, federated query, agent identity, observability

Together, these five layers establish a progression from data to action: **Data → Meaning → Trust → Reasoning → Action.** Each layer plays a distinct role in enabling reliable AI outcomes. Omitting any one of them weakens the platform's ability to support agentic AI consistently.

> **The Intelligence Data Lake is an architectural discipline rather than a commercial product. Its value lies in the capabilities it brings together, not in any single technology.**

4.1 Building The Knowledge Layer for Agentic Deployment

The Knowledge Layer is a translation layer. It sits above raw data and below the agent, transforming cryptic, tabular data into machine-readable meaning. Of the five architectural layers, it often delivers the greatest impact while requiring comparatively modest investment.

If organizations prioritize one capability, the Knowledge Layer should be at the top of the list.

The examples below illustrate how the Knowledge Layer transforms raw technical assets into information that AI agents can interpret and reason over.

From raw data to knowledge: concrete examples of how it works.

Raw data platform	Knowledge Layer
rev_adj_net_q (column)	"Adjusted Net Revenue (Quarter), excludes returns"
cust_id ↔ ord_id ↔ prod_id (FKs)	Customer purchased Product via Order
status = 'C' (stored value)	Churned: no activity 90+ days, per FY26 policy
embedding [0.42, ...] (vector)	Refund policy doc, FY26, applicable in IN region

Four Components of Knowledge Layer

1. **Semantic / Metrics Layer:** Single source of truth for measures: revenue, churn,and LTV. Tools: dbt Semantic Layer, Cube, AtScale**.**

2. **Ontology + Knowledge Graph: Entities and relationships, formally encoded. Enables multi-hop reasoning. Tools: Neo4j, Neptune, graph on lakehouse.**

3. **Vector / Embedding Layer:** Embeddings on structured and unstructured assets, including schema and business rules. Tools: pgvector, Pinecone, Weaviate.

4. B**usiness Rules and Policies:** Encoded decision logic: eligibility, SLAs, refund policies, and compliance gates.

How the four fit together: the semantic layer sits closest to the agent and defines what each measure means. The ontology and knowledge graph encode how entities relate, giving those definitions something to traverse. The vector layer indexes the same schema, documents, and rules so the agent can retrieve them by meaning rather than exact match. Business rules sit on top and constrain what the other three are allowed to conclude. Entity resolution is upstream, in the Data Foundation, supplying the resolved identities all three depend on.

One architectural note: Entity Resolution is foundational infrastructure. It belongs in the Data Foundation, not the Knowledge Layer. The Data Foundation establishes identity, while the Knowledge Layer establishes meaning.

4.1 a) Ontology In Action: A Worked Example

Question to the agent: "Why has customer growth declined in Q4?"	**Question to the agent:** "Why has customer growth declined in Q4?“
Without KL: "Growth likely slowed due to some mix of competition, seasonality, or marketing — churn and acquisition trends are unclear from current data. Recommend deeper analysis by marketing/finance."	**With KL:** "Growth fell from Acquisition ↓18% (Paid Search CTR ↓31% on 'Q4 Refresh') and Churn ↑6% (62% cite Pricing, tied to NewCo's 15% price cut), both starting in October." traceable to Channel→Campaign and Reason→Competitor paths.
Without a Knowledge Layer, answering this question requires a human analyst. the agent can't decompose the metric or trace it to root causes it just pattern-matches across disconnected reports and produces a plausible-sounding but unverifiable, often wrong, answer with no traceable path back to the data.	**With the Knowledge Layer in place,** the agent can reason through the problem.

This is what happens when there is a knowledge layer

Decompose:
The Semantic Layer defines customer growth as acquisition + expansion – churn. The agent retrieves each component.

Traverse:
Acquisition is down 18%, churn is up 6%, and expansion remains flat. The agent prunes expansion from further analysis, allowing it to focus on the most relevant factors.

Drill down via the ontology:
Acquisition → Channel (Paid Search ↓44%) → Campaign (Q4 Refresh, CTR ↓31%). Churn → Reason (Pricing, 62%) → Competitor (NewCo, -15% price launch).

Correlate and cite:
Both anomalies began in October. The agent synthesizes a response using traceable knowledge graph paths.

This example illustrates reasoning rather than simple retrieval. Every conclusion is traceable through the underlying ontology.

4.2 Designing the Intelligence Data Lake

Build incrementally:
Scope the implementation to one domain at a time. Avoid trying to solve everything at once and expand once the approach has been proven.

Make it machine queryable:
Agents interact through APIs, not user interfaces. Every access path should be exposed through a programmatic contract.

Version everything:
Knowledge evolves over time, so manage it with the same discipline as code.

Link every decision to trust signals:
Meaning should always be accompanied by confidence, lineage, and provenance.

5. Build: The Capability Roadmap For Overall Agentic AI Implementation

The roadmap comprises ten capabilities organized across three implementation phases. The sequence is intentional, with each capability building on the one before it.

The **Assess stage** establishes where you stand and identifies the capabilities that need to be strengthened. The **Architect stage** defines the target Intelligence Data Lake. The **Build stage** translates that architecture into an implementation roadmap.

Phase 1. Foundation: Identity and Integrity

- **Data Catalog and Lineage.** Discoverable metadata.
- **Data Contracts.** Machine-verifiable producer ↔ consumer agreements.
- **Entity Resolution / MDM.** Canonical Customer, Product, Transaction

Phase 2. Knowledge: Meaning and Reasoning

- **Metrics / Semantic Layer.** Start with ten key business metrics.
- **Ontology Definition.** Formal entities and relationships. Start narrow.
- **Knowledge Graph.** Operational graph store. Multi-hop enabled.
- **Vector Index Layer.** Embeddings on schema, docs, and rules.

Phase 3. Enable: Access and Governance

- **Agent Tool-Use APIs.** Structured functions agents can safely call.
- **Trust Signals and Quality SLAs.** Freshness, accuracy, confidence — exposed inline.
- **Agent Identity and Governance.** RBAC for non-human actors. Audit trails. DPDP / GDPR compliance.

How this maps to the Intelligence Data Lake:

The implementation roadmap aligns with the architectural stages, but the capabilities are introduced incrementally rather than sequentially. Phase 1 establishes the Data Foundation while introducing foundational trust capabilities. Phase 2 builds the Knowledge Layer. Phase 3 extends the Agent Interface while strengthening trust, governance, and operational capabilities.

A realistic pace is six to nine months for the first domain in large enterprises with a dedicated team and executive sponsorship. The second domain typically takes half as long. By the third, established patterns and reusable capabilities accelerate delivery.

Team profile for a first-domain build: the team is small and senior. Expect one data architect to own the semantic and ontology decisions, two data engineers for pipelines, contracts, and graph or vector stores, one platform engineer for the agent interface and tool APIs, and a part-time governance lead for agent identity and audit. A domain expert keeps the work anchored to real business questions. Executive sponsorship is the only non-negotiable prerequisite, because without it, cross-team data contracts in Phase 1 stall.

5.1 Start Here: Your First 90 Days

These are the Phase 1 activities, scoped to a single quarter.

1. **Run the Data Platform check and the AI Readiness check** on your top three data domains.
2. **Pick one business entity** and build its canonical definition end-to-end.
3. **Deploy a metrics layer on one critical metric,** such as revenue or churn, and establish a single, consistent definition.
4. **Write your first data contract** between a critical producer and consumer.

5.2 Five Mistakes To Avoid

1. **Having a data lake does not mean being AI-ready.** Data volume alone is not a measure of AI readiness.
2. **Skipping the Semantic Layer.** It is often rebuilt later under greater time and delivery pressure.
3. **Treating agents like human users.** Agentic systems require different access patterns, latency expectations, and identity models.
4. **Operating without data contracts.** Upstream changes can silently disrupt downstream agents.
5. **Treating Governance as an afterthought.** Retrofitting compliance later is significantly more costly than designing for it from the outset.

6. How Tiger Analytics Partners On The Journey

Most enterprises can begin this journey with their own teams. Where they typically seek support is in the decisions that are difficult and expensive to reverse. Every enterprise we partner with on this journey moves through four phases, with the starting point determined by its current level of maturity. At each stage, our objective is to build your team's capabilities, not create long-term dependency.

Phase	What we do	What you get
1. Assess	Run the AI Readiness Audit and Data Health Check across your priority domains.	A scored diagnostic, gap-map, and a prioritised backlog aligned to business value.
2. Architect	Co-design the Intelligence Data Lake for your context — Optimize path or Leapfrog path.	A reference architecture, data contract templates, and an ontology blueprint.
3. Build	Implement the Knowledge Layer for one domain, with your engineers as embedded co-builders.	A production-grade Knowledge Layer, agent tool-use APIs, and trust-signal instrumentation.
4. Run	Operationalise agent identity, governance, and observability. Enable the second domain.	AgentOps runbooks and a repeatable pattern for domain 2+.

What differentiates our approach

1. **Data engineering first, AI second.** We start with the foundation because that is where the failures live.
2. **One domain, done right.** We resist the temptation to boil the ocean. Depth in one domain teaches you the rest.
3. **Embedded, not delivered.** Your engineers build alongside ours. We leave capability behind, not just documentation.
4. **Regulator-ready by default.** Especially for BFSI and healthcare — DPDP, GDPR, sectoral regulators built in from day one.

> **Where we add most value:** Clients consistently tell us that our greatest value comes at two points in the journey. The first is during the semantic-layer decisions in Phase 2, where architectural choices are difficult and expensive to reverse. The second is during the operating-model transition in Phase 4, when organizations move from project delivery to platform capability and require new roles, governance models, and success metrics. For organizations at the beginning of their AI journey, these are the conversations we prioritize.

7. The Reframe

If you remember one thing from this paper, remember this single question.

> ***The question is not how to apply AI to your data, but whether your data is structured for AI to act on it reliably.***

That reframe shifts the conversation from the visible part of the system (the model, the prompts, and the agents) to the data foundation where the greatest leverage exists. It places accountability where it belongs, on the data foundation.

The Intelligence Data Lake is an architectural discipline rather than a vendor product. Most of the capabilities required already exist within modern data engineering teams. What has been missing is the framing: recognition that the consumer of data has changed and that the platform must evolve with it.

7.1 Four Takeaways

- **The AI Readiness Gap is structural.** It is a data foundation problem rather than a model problem, which means organizations can address it through deliberate architectural choices.
- **The Knowledge Layer is the most critical missing capability,** bringing together semantics, ontology, vectors, and business rules, on a foundation of resolved entities and trusted data.
- **Two paths converge at Level 4.** Optimize if you have legacy investments. Leapfrog if you are greenfield.
- **Start with one domain, one entity, one data contract.** Prioritize depth before breadth. The benefits compound over time.

> *The platforms that succeed in the agentic era will not necessarily be those with the most data. They will be the ones whose data is most worthy of the AI built on top of it.*

About the author

Technology Leader with 25 years of IT experience spanning Data Engineering, Generative AI, Machine Learning, Big Data, Cloud, Data Warehousing, Business Intelligence (BI), and Analytics systems. Muthu Govindarajan has successfully played diverse roles in technology leadership, practice leadership, product ownership, and architecture, consistently driving innovation, scaling practices, and delivering enterprise-wide digital transformation.

Muthu Govindarajan
Partner - Engineering
Tiger Analytics