HEXAWARE

Infinite Trust – Data Readiness for AI

Girish Pai
Executive Vice President & Global Head – Data, Hexaware

Agentic AI across the Data & Analytics Value Chain

HEXAWARE

1

Data Readiness for AI

2

Agentic Data Engineering

Dramatic efficiency improvement and speed to market through agentic automation of the data supply chain

Delivers Infinite Trust by way of accurate, reliable, complete, governed, secure contextual data for advanced AI solutions at enterprise scale

3

AI Solutions

Turns curated datasets into instant business value through autonomous & conversational analytics

“Vibe Analytics”

Without trust in data, intelligence never leaves the lab

Your agents are only as smart as your data's meaning — and only as trusted as the governance behind it

60%

of AI projects will be abandoned through 2026 — undone not by a lack of data but by data that isn't either available, standardized, contextualized, trustworthy or timely.

Current projected size of the AI readiness for Data services expected to triple by 2032

$40B

HEXAWARE

Why Trust Breaks Down–
The Problems Aren't New, Agentic AI Just Raises the Stakes

Every generation of AI has the same data fragmentation, meaning, trust and governance gaps — agentic AI brings these to the forefront.

Knowledge Split: Facts vs. Context 01

Facts exist in databases. Context resides in policies, memos and PDFs.

Manual Assembly Cannot Scale or Achieve Coherence 02

Connecting a number to the document that explains it is still a manual research task today

Data Is Not Available at Agent Speed 03

Different data sits in different systems, owned by different teams, Pipelines still take hours or days

Human-Ready Is Not Agent-Ready - Intent Doesn't Transfer 04

Humans intuitively understand insights & intent from dashboards. Agents need explicit structure and rules to infer meaning

Policy Is Not Runtime 05

Role, risk and decision rights are documented — not enforced live.

Trustable data essential for reliable AI 06

The impact of errors from bad data are exponentially amplified in an AI world

The gap isn't more data. It's connecting meaning, trust , governed authority and accountability at the speed and precision agents require.

The Trust Gap Is Real – Closing It Is Harder Still

HEXAWARE

Most enterprises don't lack ambition for agentic AI — they lack a fast, affordable, durable way to close the trust gap before priorities shift again.

Timing Trap

Data-readiness programs take 12–18 months. Data, models , regulations move faster- by the time they finish, the target has moved

Front-Loaded Cost

Large consulting engagements front-load cost and diagnostics long before any AI use case shows business value.

Talent Shortage

Semantic, governance and agent-engineering skills are scarce, expensive and hard to build in-house fast.

01 02 03 04 05 06

Too Many Priorities

Data, security, compliance and business teams compete for the same budget, people and time

Compliance Bottleneck

Every new data path must clear security, privacy and compliance reviews before it hits production.

No Single Owner

Data, IT, security and business each own a piece of the problem —no one owns all of it.

Overcoming these with a traditional approach can cost tens of millions and take 18–24 months

HEXAWARE

Our Approach: AI-Led, Human-Enabled Readiness Assessment

A fast, parser-based way to assess current state deterministically, clarify priority outcomes and define practical path to value with an enterprise intelligence blueprint

Inputs / Discovery

- Current architecture
- Data landscape
- Data availability & access
- Data quality and trust metrics
- Semantic context & knowledge
- Use cases & business outcomes

AI-Led, Human-Enabled Assessment Approach

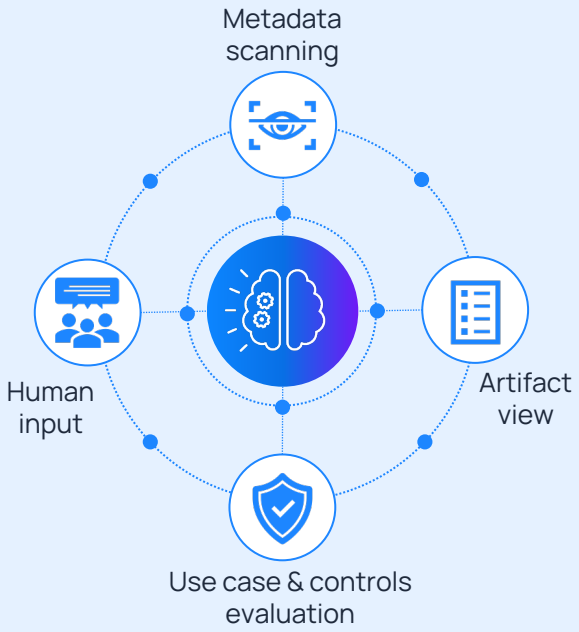

Enterprise Intelligence Blueprint

Semantic Knowledge Layer
Context enriched, curated data

Model & Agent Intelligence Layer
Governed intelligence to power AI use cases

Built in **6-8 weeks**

..on your **existing investments**, not around them

Deliver the blueprint in iterative sprints of 6–8 weeks for each functional area for about $1M.

HEXAWARE

Making Data AI-Ready

Making Data AI-Ready

Case Studies

HEXAWARE

AI-ready Data: Transforming Investment Data Intelligence for a large 'Investment Pension & Insurance' provider

The Challenge

- The client lacked portfolio transparency resulting in **latency in the Data-to-Decision cycle**
- 10+ siloed data sources resulting in **40-50% of** analyst **time** spent in **manual data analysis**
- **Over-reliance on external 3rd parties** (£2-3 Mn spend on verification queries with 48 hr delay in receiving position breach data).

Our Solution

- Built a cloud-native **investment data platform on GCP, engineered for scale, governance, and reuse** across business decisions.
- Developed **AI-ready investment data products** covering Fund Hierarchy, Fund positions, Security Master, Monthly performance and ESG attributes
- Enabled trusted, decision-ready; Chief Investment Officer executive dashboards.
- Built an **enterprise semantic knowledge layer** (taxonomy, business semantics, ontology, knowledge graph) to standardize meaning & deliver trusted, contextual, decision-ready data.
- Enabled **self-service natural-language querying in Vibe Analytics**, allowing Investment Desktop users to interrogate data directly.

Business Outcomes

70%
Reduction in manual data preparation time, improving data coherence.

85%
Faster time-to-insight for market event impact assessments.

£2m
Annually saved in 3rd party data request spend & other vendor servicing costs.

AI Ready Data: Modernize to Governed AI-ready Data platform for a Big 4 Audit Firm

The Challenge

- Legacy SQL Server EDW **limits next-generation analytics** across Finance, Procurement, Operational IT and Clients & Markets.
- High on-premises SQL Server **maintenance costs and infrastructure dependency**.
- **Limited scalability** for growing analytics workloads

Our Solution

- **Agentic discovery** of the existing landscape on Hexaware's AMAZE platform, producing a complexity report, object inventory, dependency & relationships between objects, source code comprehension, etc.
- **Automated conversion & migration** of on-premises EDW procedural logic to Databricks notebooks, with orchestration aligned to Databricks **lineage and workflow dependencies**.
- Established a factory model roadmap with parallel migrations for EDW and reporting.
- **Built enterprise semantic model and business semantic layer** to enable self service in the target data platform
- **Introduced Databricks Unity Catalog for centralised governance**

Business Outcomes

5,000+
Stored procedures & SSIS packages migrated to Azure Databricks notebooks.

8,000+
Tables and views moved to Delta tables across Bronze, Silver and Gold layers.

40-60%
Automated discovery and conversion through the AMAZE platform.