

AI for Business

Shantanu Baruah

President & Global Head – Healthcare, Life Sciences &
Insurance and Head – North America (Hunting), Hexaware

AI for Business Paradigms

We continue to build capabilities in the following areas to execute our strategy.

01



Solving the Hard Problem

02



Entering New Markets



Cases in Point



For a **Top 4 Audit Firm**, building agents to maximize sales opportunities by bringing right stakeholders, and knowledge context in one pane—anticipating **75% higher pipeline and 25% more closure**



For a physician-owned mutual insurer dominant in medical malpractice, implemented an agentic AI solution to read thousands of pages of documents to reduce **MFQ cycle time from 60 days to less than 5 days and improve accuracy by 25%**

From Protocol to Submission, on One Source of Truth

PROBLEM STATEMENT

 **Every clinical study is still rebuilt by hand**

- Protocols are re-keyed into EDC, IXRS, CTMS, eCOA, and lab systems over many weeks
- Each manual translation introduces errors, delay, and version drift
- There is no single source of truth from protocol through to submission

HEXAWARE SOLUTION

Author once as USDM, publish to one hub

-  **Machine-readable USDM 4.0 core**
Protocol parsed once into structured CDISC data, validated
-  **One SDR hub, every consumer pulls**
Documents and clinical systems auto-configured via API, one source
-  **Agent-native and GxP-traceable**
Confidence-scored, human-reviewed, traceable to source, on AWS

QUANTIFIED BENEFITS

30 to 50%

Faster study startup, the DDF industry target

9 Domains

CDISC entities extracted per protocol vs manual re-keying

4 Plans

L2 document set auto-drafted, 12 files in parallel

1 Source

Canonical USDM record, the EDC build cycle removed

HOW IT WORKS

1 Ingest

Protocol PDF or DOCX landed in the secure console

2 Extract

9 CDISC domains parsed, terminology mapped

3 Validate

USDM schema checks, confidence scoring, human review

4 Publish

Assembled USDM synced to the SDR, documents drafted

Unlocking the Clinical Value Trapped in CCDs

PROBLEM STATEMENT

 **70% of clinical detail lives in CCD narrative, invisible to claims-based programs**

- Diagnoses, severity, and status are buried in free text, never reaching a claim
- Each CCD arrives in a different template, section order, and vocabulary across HIE and EHR sources
- Manual chart abstraction is the only fallback: slow, expensive, and impossible at population scale

HEXAWARE SOLUTION

CCD-native clinical NLP, grounded and confidence-scored

-  **400+ entity extraction**
Diagnoses, meds, procedures, and results pulled across every CCD template
-  **Mapped to standard terminology**
Auto-coded to SNOMED-CT, ICD-10, RxNorm, and LOINC
-  **Clinical reasoning engine**
Reconciles contradictory CCD, HER, and claims by recency and source

QUANTIFIED BENEFITS

400+

Clinical entities extracted per CCD vs ~15 discrete claim fields

3 to 4x

More documented conditions surfaced for risk and quality vs claims alone

99%+

De-identification accuracy, validated against expert review

< 24 hrs

From CCD receipt to coded, queryable record vs weeks of abstraction

HOW IT WORKS

1 Ingest

CCDs pulled from HIE and EHR via FHIR and C-CDA, landed as-is

2 Extract

400+ entity NLP parses narrative, sections and context

3 Normalize

Mapped to SNOMED, ICD, RxNorm, LOINC with confidence scores

4 Publish


Written to the longitudinal record with full source provenance

AI for Business

Customer Showcases

Eravi Gopan

President & Global Head – Technology,
Products & Platforms, Hexaware



AI for the business of AI. Net New TAM for Hexaware Services. A 2T + economy.



FOUNDRY LABS & CONSUMER
ELECTRONICS

Picks and Shovels for the AI economy

Making AI work for all of **HUMANITY**

Making AI work for all **ENTERPRISES**

Making AI **COMMERCIALLY SUCCESSFUL**



AI INFRA PROVIDERS -
STORAGE, COMPUTE, MEMORY,
SEMI, DCS, HI-TECH MFG


AI workforce for unlocking capacity

Compute, Memory, expertise &
time constraints

MAXIMIZE the business **OPPORTUNITY**

MINIMIZE the **FRICTION** to business

Making AI models work in the real world



2026:
Hexaware's
RLHF win with
a Top AI lab

Hexaware's **RLHF-as-a-Service** is helping the AI Lab expand to global markets

Hexaware **Pods** of scholars and linguist, process experts and quantitative methods provide feedbacks on cultural, and domain nuances



\$27 B near-term services
Target Addressable Market (TAM)



opportunity for AI services across
Pretraining and Post training value chain as
enterprises move to advanced alignment to
real-world understanding

LLMs/platforms | AI co-pilots/Agents | LLMs & AI applications | AI devices | Multimodal channel for models



Hexaware's Product Ownership of AI workforce strategy - Capacity unlocked becomes capital deployed - for rapid growth into the company valuation

ZERO MEETINGS Non-IT fab roles (engineers, operators, process owners, site leads) Extract and agent-coordinate shared information-handling and coordination work across many roles	ZERO TICKETS Reactive to proactive - L1 & L2 IT support functional units Rebuild entire IT support functions end-to-end as agentic	EMBEDDED AI SCOUTS 7 scouts embedded across business units Learn the functional unit deeply (60%), then identify and build the case for agentic shift (40%)
Operational gain — meetings become the exception, not the norm	Self-funding mechanism — freed capacity redeployed into the AI workforce execution	Releases capacity + serves as the discovery/sensing engine for grounds up AI transformation



HEXWARE

EXCLUSIVE AI SOLUTIONS PARTNER


Fortune 500 Hi-Tech Major & The AI CapEx Supercycle

Market Data Offering

Ravi Vaidyanathan


President & Global Head – Financial Services, Hexaware

The Market Opportunity



! *Clients are deeply concerned about the cost of market data—and the lack of innovation.*

Hexaware Play



The Hexaware Play

INITIAL OFFERING



Optimize Market Data
Spend for customers

Derive broader value for client
investments in Market Data

- Replace / Improve selected products subscribed by clients
- Create AI-native data sets that don't exist today
- Augment Static Data sets with AI driven analytics
- An AI-driven platform to manage contracts, usage, attribution, and effectiveness of market data across the entire enterprise

Leading to



EXPERIMENT BEING A MARKET DATA PROVIDER



Standard AI-Native Datasets

Static Datasets Augmented with
AI driven insights

Custom Data Products

- ARR revenue stream
- Underserved niche
- Alternate and Private Market Data sets
- Proprietary indexes, analytics and insights

TAM

Initial Offering

Market Data
Spend
Optimization

\$50B TAM

ACQUISITION & LICENSES

20% - 30% Cost
OptimizationAI-driven contract, usage,
and attribution management

DATA OPERATIONS

50% Cost
OptimizationMost labor intensive-biggest
automation lever

PRODUCT & ENGINEERING

70% Improvement in
Time to MarketCloud-native pipelines and
API-driven product creation

BUSINESS CASE

\$10M Per Client;
Target 5 Per YearGain-share mechanism with
a Y1 minimum target of 20%
cost savings

As a Market Data Provider

Custom
Products

\$5B - \$10B TAM

CORE SCOPE

Custom Datasets

Proprietary indices,
customized factor
formulations, and bespoke
data feeds

GROWTH DRIVER

~10-20% of Industry

Fast-growing, client-
specific APIs and bespoke
quantitative datasets

EXCLUSIVITY MOAT

Avg. \$80k ARR

For 20 client firms per asset
type

BUSINESS CASE

\$1.6M ARR Per
DatasetTargeting 4 dataset
launches in a yearAI Augmented
Data

\$8B TAM

REAL-TIME TRADING

\$17B - \$20B Segment

Active trading desk feeds
and ultra-low latency real-
time data

DESKTOP TERMINALS

\$10B - \$12B Segment

Consolidated screen platforms led
by institutional terminal habits

REFERENCE & PRIVATE

\$13B - \$18B Segment

Fastest growing niche indexes and
private illiquid assets

BUSINESS CASE FOR MacroIQ

\$6M for 10k Users

*Pricing subscription @ \$50
per user per month.
Planning to launch 4
products in a year*

Market Data Cost Optimization Strategy – An Illustration

Strategic Blueprint for \$1T AUM Asset Manager | Ingestion-to-Insight AI Transformation

Current Spend \$100M / yr	Target Savings \$28M - \$35M	Efficiency ROI 3.2x
------------------------------	---------------------------------	------------------------

 Current Spend Landscape Baseline \$100M Data Footprint	 Agentic AI Replacement Automating Scraping & Structuring	 AI Duplicate Elimination Overlap & Seat Consolidation	 Alternate Data Sources Modern Cloud & Open APIs
Financial Terminals \$42M (42%) Indices & Benchmarks \$28M (28%) Reference & Alt Data \$30M (30%)	SEC & Regulatory Filings -\$8M Savings Macro & Central Bank Data -\$4M Savings ESG & Transcripts Parsing -\$5M Savings	Terminal Seat Rationalization -\$9M Savings Enterprise Index Centralization -\$4M Savings Corporate Action Normalization -\$3M Savings	Cloud Data Marketplaces High Value Open Financial Micro-APIs Lower unit Cost Synthetic Benchmark Engines Custom Blends



Strategic Shift: "Data Ingestion" → "Direct AI Insight Generation"

Bypassing raw data aggregation pipelines by using LLMs to ingest primary sources directly into alpha & risk decisions.



Market Data Spend Optimization




Contract, Entitlement, and Usage Control

Every license, permission and actual usage in one place—so you can see and challenge what you pay for



AI-Automated Data Operations

Data handling is automated, with fewer people



Cloud-Native Delivery

Marketplace and cloud distribution—new datasets onboarded in weeks, without an integration project



Managed Data Products

AI-enabled data pipelines, quality automation and observability run as a managed service for institutional clients



Built on our Market Data Management System blueprint—an existing, documented design covering contracts, usage, attribution, and effectiveness across the enterprise

AI Augmented Data

Focus Area Opportunity and Example Datasets · AI Augmentation of Standard Data Sets



Reference and Pricing Data

The mission-critical “plumbing”—sticky, recurring, and worth more as illiquid assets need evaluated pricing

Examples: AI augmentation provided to add commentary on reference, price data provided by the external vendors



Analytics and Workflow

Where raw data becomes decisions—the premium sits with analytics and AI copilots embedded in the workflow

Examples: AI explained performance attribution, performance attribution, fixed-income analytics (OAS, duration), ESG scores, index analytics



Alternative Data

The differentiated-alpha frontier; the biggest unmet need is making it clean, linked, and AI-ready

Examples: Web-scraped pricing, NLP news sentiment, shipping lines (AIS), localized / seasonal trading patterns



Private Markets Data

Fastest-growing pocket—low transparency in PE / VC / private credit means buyers pay a premium for coverage

Examples: Private company financials and valuations, fund returns (IRR / TVPI / DPI), deal databases (M&A, PE / VC rounds), cap-table & LP / GP data

Source: 2025 industry estimates (Mondo Visione, Traders Magazine, Investopedia, provider financials). Coral = focus areas marked in source. Ranges are estimation bands.

Custom Data Product

Focus Products and Example Datasets · Buyers: Asset Managers · Investment Banks · Wealth Platforms · Managed Services



Regulatory Datasets

Compliance-driven

Non-discretionary spend that grows every time a new rule lands—budget is mandated, not optional

Examples: KYC / AML data, Sanctions and PEP lists, Basel / CRR risk inputs, LEI / taxonomies



ESG Datasets

High growth

Regulation (CSRD, SFDR, EU Taxonomy) is forcing demand while persistent data gaps keep prices high

Examples: Scope 1 / 2 / 3 emissions, climate physical and transition risk, controversy screening, board & diversity governance, taxonomy alignment



Private Market Databases

Bespoke coverage

Coverage the big terminals don't reach—productized private intelligence sold as a database

Examples: Proprietary company and deal databases, fund-benchmark custom cuts, private-credit facility data, sector-specific private datasets



Client-Specific Data Feeds

Stickiest revenue

The most defensible revenue of all—bespoke, managed, and expensive for the client to switch away from

Examples: Bespoke golden-source feeds from niche leaders, custom factor models, white-labelled indices, portfolio-specific reference and pricing

Source: 2025 industry estimates (Mondo Visione, Traders Magazine, Investopedia, provider financials). Coral = focus areas marked in source. Ranges are estimation bands.

MacroIQ

Every firm buys from the same vendors, reads the same releases, and reaches the same conclusions, days apart. Advantage no longer comes from owning data. It comes from how fast that data becomes a decision.



MARKET NEED

Trusted Data

The signal is scattered

Official prints sit across BLS, BEA, the Fed, Treasury, IMF, and OECD; forward-looking views are buried in hundreds of institutional publications. Consolidating who forecasts what—with revisions and citations—is manual work today.

Research teams face overload

Analysts reconcile actuals, estimates, and revisions by hand across formats and vintages. The questions they actually ask—“what did consensus believe then?” “which source is reliable?”—have no off-the-shelf answer.

The bar is enterprise-grade trust

Institutional buyers demand provenance on every number, measured accuracy, freshness as a service level, and proof their proprietary data cannot leak. Vendors are judged by what their platforms visibly enforce.



MacroIQ


Standard macroeconomic data, augmented with AI-driven, analyst-vetted insight—data and intelligence in one package



FOUNDATION

The Data

- **US macro indicators**
Historical actuals plus estimated values
- **Multi-source forecasts**
Governments, economists, analysts & financial institutions
- **Article-sourced**
Drawn from published research and articles



WHAT MAKES IT DIFFERENT

The Intelligence Layer

- **Summary + citation**
5–7 lines on each indicator's estimate and why it was revised
- **Persona-specific takes**
Summaries tailored to each type of user
- **AI impact analysis**
What the move means, generated automatically
- **One clear visual**
Data and summary combined in an easy-to-read chart



BUILT FOR

The Users



Asset managers & research



Banking treasury / ALM



Insurance underwriting & risk



Economics & research teams



Academics & students

AI-generated, human-analyst vetted. Unlike typical marketplace datasets — not just data, but the intelligence to act on it.

AI in Professional Services

Kush Gupta

Senior Vice President & Global Head – Professional Services,
Hexaware



Creating non-Linearity in Professional Services



SCALE



> 50%

AI-influenced Revenue

75B+

Tokens consumed per month for AI dev

250M+

Transaction / day for AI platform co-developed

150+

Agents in production



EXECUTION



50%

Outcome / Fixed price projects

100+

AI projects

70%

AI projects use Kanban

90%

AI-ready workforce



PARTNERSHIP



databricks

#1 partner for Databricks in PS



Microsoft

#1 partner for Big 4



FACTORY

Largest certified partner in Big 4 ecosystem

servicenow

Large Dev partner for Big 4

IMPACT & OUTCOMES



Trusted tech partner for **3 of the Big 4 firms** behind their audit and tax products



10x higher coverage of Audit and 35% reduced timeline leading to **\$5M annual savings**



Designed and developed the **tax filing platform** leading to 40% reduction in tax filing



AI-enabled eDiscovery platform with **Legal SLM** for **77% higher precision** and less than 5% Manual work

Impact of AI-Enabled Audits

We are optimizing the way audits are carried out.

AI-Enabled Audit Automation

20+ AI micro-products developed to replace Excel-based processes across the audit practice

\$9M annual savings—nearly self-funded

Agentic Audit Framework

Multi-agent framework untangling platform dependencies across the audit organization

150k auditors reached—unified task assignment, guidance, and review

Audit Analytics Platform

Digital Audit Platform embedding analytics across the entire audit lifecycle

Full data coverage instead of random sampling

AUDIT VALUE CHAIN TRANSFORMATION					
	PLAN (Intake, scope, and planning)	>> UNDERSTAND (Research and content mgmt.)	>> EXECUTE (Execution)	>> VALIDATE (Substantive procedures)	>> CONCLUDE AND REPORT (Audit analytics)
FROM	Manual effort for conflict checks	Policy review, scattered knowledge	Sample testing, periodic assessments	AR / AP confirmations, excel balancing	Manual review, validation via judgment
TO	Agent-based scoping, risk scoring, and conflict monitoring	Automated document analysis and contract	Full-coverage testing, financial reconciliation, and anomaly detection	AI workpapers with autonomous AR / AP confirmation	Insight-driven documentation with auto report generation



Intelligent Tax Journey with AI

From a periodic, sampling-based compliance cost center to a continuous, full-population, advisory-led tax function

WE HAVE EXECUTED FOR LEADING GLOBAL FIRMS

Automated Framework for Mobility Platform



Intelligent AI-driven automation framework built on the Factory AI platform

60-70% effort reduction and 100% workflow traceability

AI-Powered Tax Filing Automation



Automated tax return preparation using AI-driven document intelligence

70% faster questionnaire population

AI-Assisted QA for Tax Portfolio



Scaled AI testing across tax applications with 60k+ test cases per release

60-70% effort reduction in test design and defects

TAX PROCESS TRANSFORMATION (Manual → AI-enabled)

	Data collection and mgmt.	Determination and provision	Compliance and return prep	Review and validation	Planning and controversy
FROM	Manual data pull, fragmented data	Manual provision and account calculations	Regional return prep and filing	Sample-based workpaper review	Reactive planning and labor-intensive research
TO	Automated extraction and classification of source data	Automated provision and real-time computation	Auto-generated returns and multi-jurisdiction e-filing	Anomaly detection and continuous validation with HITL	Real-time planning with AI research and controversy support

THE IMPACT



~17 to ~8 wk

Compliance cycle — with HITL



~240 hours

Saved / year / professional



Up to 70%

Less time per return



~\$19k

Annual value saved per employee

Reinventing eDiscovery With AI

The same legal process, re-engineered. AI compresses cost, time, and risk at every stage of the discovery lifecycle

OUR SUCCESS STORIES IN LEGAL

Enterprise Data Platform

Intelligent, insights-friendly enterprise data platform creation

50% reduction in platform enablement
—1.5x faster benefit realization



AI-Enabled eDiscovery Platform

Built a customized AI-enabled eDiscovery platform

10x–20x return on software investment



Legal Copilot & Agentic Solution

Contract intelligence, AI-based case intake and matter onboarding

AI copilots across the legal workflow



eDISCOVERY LIFECYCLE TRANSFORMATION

	Identify	Preserve and collect	Process	Case assessment	Review	Produce
FROM	Manual interviews and data-source guesswork	Over-collection and manual legal holds tracking	Slow de-duplication and indexing	Late, gut-feel scope and cost estimates	Page-by-page keyword review	Manual redaction and Bates numbering
	⌵	⌵	⌵	⌵	⌵	⌵
TO	Data source discovery, legal holds recommendation	Collection planning; chain-of-custody monitoring	De-dupe, OCR, metadata enrichment, sensitive info	Case intelligence, volume, cost and theme projections	Relevance and privilege review, contract and comm. summaries	Automated redactions, QC sampling

IMPACT

Hours, not weeks

case assessment and summarization at scale

80% decrease

in manual review volume

40–70%

lower review costs with AI—accelerated case timelines

77% precision

outperforms first-pass human review by up to 36%

Only ~5%

of documents need human review—vs 25–30%

We Are Building the AI Layer for Professional Services.

We invite you to build it with us.

