The Emerging Autonomous Enterprise: How Leaders Are Redefining Enterprise Value with Modern AI

teradata.

2026: The autonomous enterprise leap

An autonomous enterprise uses its knowledge, applies agentic reasoning, and delivers faster, cost-effective, and scalable outcome for its customers.

In 2026, autonomous enterprise is a business reality, not a future state. By integrating enterprise knowledge (context) with agentic execution that can reason, decide, and act, enterprises move beyond isolated task automation. Autonomous operations now run end-to-end business processes across finance, supply chain, and human resources. Industries including logistics, retail, healthcare, defense, and manufacturing stand to benefit the most, as autonomy rewrites the limits of scale, speed, and performance.

Enterprise value creation is no longer theoretical. Banking can now scale personalized financial guidance to every customer. Healthcare can materially improve access and first-call resolution by up to 30%. Telecom can compress service issue resolution from hours to minutes. With the right technology foundation,the autonomous enterprise delivers measurable business outcomes.

This is the tipping point in which top-quartile enterprises operationalize autonomy, creating a permanent performance gap that traditional operating models cannot close.

Modern AI powers enterprise autonomy

Predictive AI is limited to estimating outcomes and informing human decision-makers. Autonomous enterprises require AI systems that ingest real-time signals, interpret enterprise context, decide across constraints, and execute end-to-end processes. Enterprise autonomy requires moving beyond predictive AI to execution-ready systems.

Bridging this gap requires more than better models, it requires a fundamentally different system architecture—modern AI.

What makes modern AI essential for enterprise autonomy:

- **Unified execution architecture:** Modern AI is a unified execution system that brings together multiple AI capabilities—predictive, generative, agentic, multi-modal, real-time, and optimization—to drive outcomes, not just insights. By operating directly inside live business processes, it turns data, models, and agents into measurable results instead of isolated predictions.
- **Runtime business context:** Modern AI operates with full enterprise context at execution time—combining real-time signals, operational constraints, and institutional knowledge—not just training data.
- **Closed-loop learning:** Modern AI learns from the outcomes of its own actions,continuously adapting behavior as conditions change.
- **Built-in governance:** Modern AI consistently operates under cost, policy, security, and governance constraints.

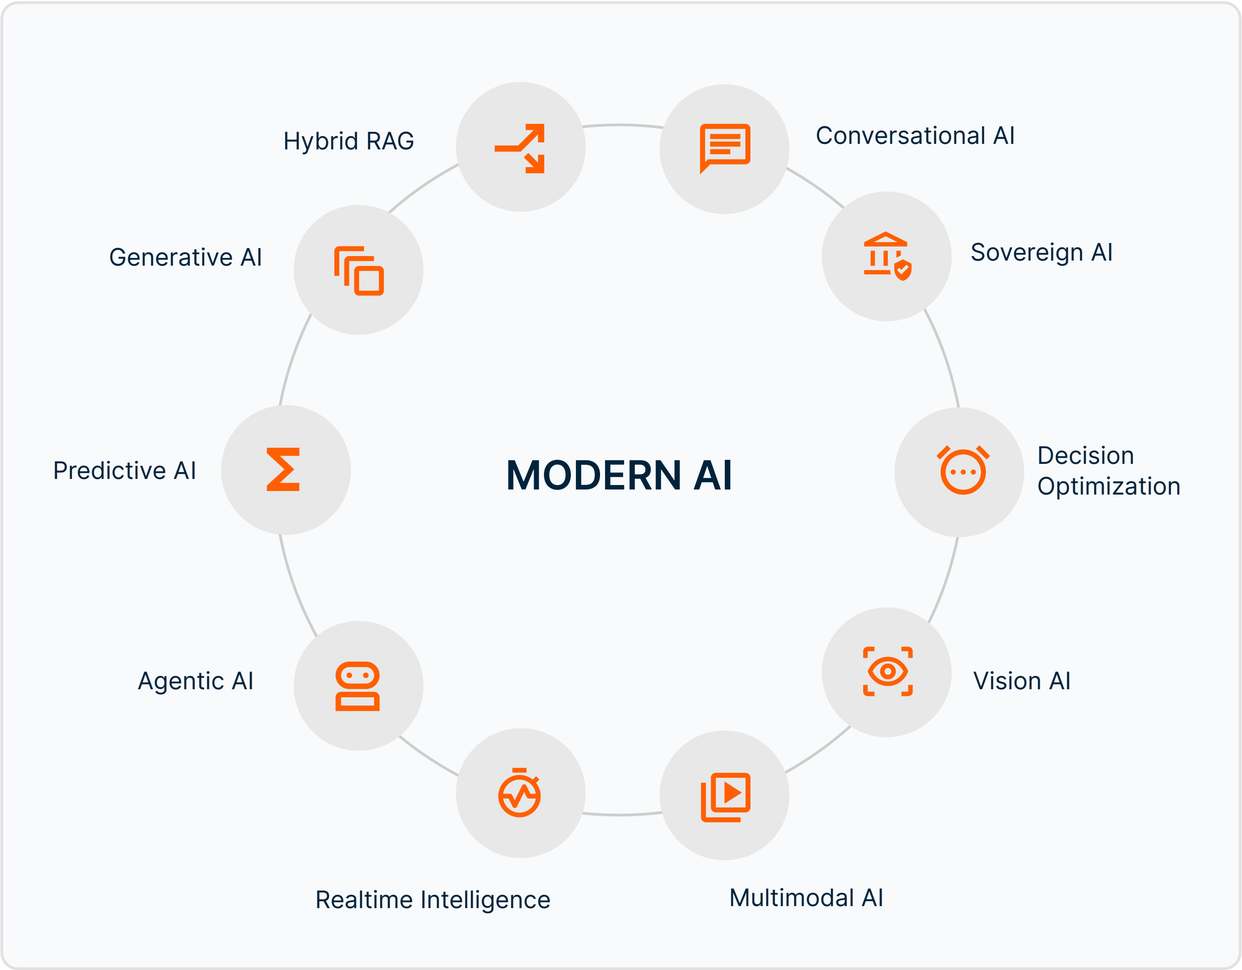

Predictive AI vs. modern AI by industry

Industry	Use Case	Predictive AI	Modern AI
Retail	Order management	Forecasts demand, stockouts, and delivery delays to generate replenishment and fulfillment recommendations.	Orchestrates realtime fulfillment with node selection (store), substitutions, order routing, and customer messaging.
Banking	Decision making	Scores credit and fraud risk to generate ‘approve or reject’ recommendations and thresholdbased alerts.	Resolves credit/fraud decisions in real time, generates reason codes and messaging, and executes followups.
Logistics & ports	Operations execution	Predicts congestion risk and delays, generating alerts and optimization recommendations for human operators.	Orchestrates berth scheduling, yard allocation, labor assignment, and vessel rerouting in real time.
Defense	Situation analysis	Analyzes historical intelligence data to identify patterns, flag anomalies, and support rootcause analysis.	Integrates sensor feeds, imagery, and operational reports to deliver realtime situational threat assessments.
Manufacturing	Continuous optimization	Forecasts yield, identifies bottlenecks, and recommends process improvements based on historical data.	Implements closed loop optimization by adjusting process parameters based on live production outcomes.
Telecom	Monitoring & reporting	Uses historical data to generate performance dashboards, thresholdbased alerts, and SLA breach reports.	Monitors network telemetry to detect service degradation and trigger remediation with followup actions.
Healthcare	Service delivery	Predicts patient risk, capacity constraints, and service gaps to generate alerts and care recommendations.	Coordinates realtime service delivery with scheduling, crossteam task routing, and adaptive care plans.

Modern AI outcomes by industry

Modern AI outcomes harness multiple AI capabilities to drive measurable ROI across industry

- Predictive AI
- Generative AI
- Agentic AI
- Multi-modal AI
- Vision AI
- Conversational AI
- Hybrid RAG
- Real-time Intelligence
- Decision Optimization
- Sovereign AI

Examples of modern AI outcomes

Defense intelligence

Visual intelligence driven protection of $15 million defense assets

Agentic logistic

Agentic AI for port congetion solution with potential cost savings of $120 billion per year

R&D AI assistant

Knowledge intelligence unifies vehicle telemetry with vectorized vehicle specifications

AI agent for every customer

A personal financial advice agent to every banking customer

Agentic hyper personalized healthcare

An agent for every patient, monitoring health and fitness data predicting likely medical issues

Maritime AI surveillance

Geo-spatial analysis of high frequency vessel movement data

Modern AI outcomes by function

Modern AI drives modern AI outcomes by operating as an execution system, not just an analytical layer. It ingests real-time signals, applies enterprise context, and executes decisions end-to-end through agentic workflows. This execution architecture enables measurable outcomes across critical enterprise functions. These are outcomes that predictive AI—limited to forecasts and recommendations—can't deliver.

Operations

- Higher operational efficiency (OEE)
- Reduced operational downtime
- Lower cost-to-serve

Customer Experience

- Reduced time to resolution (MTTR)
- Higher customer lifetime value
- Reduced customer churn

Sales

- Higher win rates
- Shorter sales cycles
- Increased pipeline velocity

Finance

- Faster financial close
- Reduced DSO
- Reduced compliance risk

Human Resources

- Reduced employee Attrition
- Improved employee engagement
- Faster time to Productivity

Marketing

- Higher conversion rates
- Lower cost per lead (CPL)
- Lower customer acquisition Cost (CAC)

R&D / Innovation

- Faster product time-to-market
- Reduced R&D cycle time
- Higher innovation throughput

Customer Experience

- Reduced fraud loss
- Improved regulatory compliance
- Lower security incidents

Vision vs. reality

Modern AI makes enterprise autonomy technically possible. Yet in practice, most enterprises remain far from autonomous.

AI still operates at the edges of the business, isolated from core systems, live workflows, and governance. As a result, the vision of enterprise autonomy is clear, but the operational reality falls short.

95%

of enterprise AI initiatives deliver zero measurable P&L impact; only 5% reach workflow-integrated execution at scale”

MIT Project NANDA: State of AI in Business **2025**

74%

of enterprises plan to deploy autonomous AI agents within two years, yet only 21% have mature governance to support autonomous execution”

Deloitte: State of AIin the Enterprise **2026**

5%

of companies generate AI value at scale, while 60% see little to no value despite sustained enterprise investment.”

BCG: The Widening AI Value Gap **2025**

Barriers to enterprise autonomy

The gap between AI capability and enterprise autonomy is not a failure of models, but a failure of system design. Today's enterprise AIenvironments are constrained by three fundamental technological barriers that prevent AI from operating as a governed, end-to-end executionsystem at scale.

Barrier 1: Fragmented systems

Fragmented enterprise AI systems separate data, models, agents, and applications. This prevents AI from executing decisions inside live business processes and forces data and models to move between systems.

Result:
Silos, integration friction, and governance risk

Barrier 2: Missing business context

Enterprise AI systems lack runtime business context, operating without trusted data, institutional knowledge, and real-time operational signals.

Result:
Context gaps, execution failure, low trust

Barrier 3: Locked-in systems

Enterprise AI systems run on closed, proprietary platforms that restrict interoperability across models, tools, and runtimes.

Result:
Execution-time lock-in, limited scalability, and increasing cost

The Opportunity Triad

In 2026, agentic AI, enterprise grade context, and hybrid interoperability converge to enable enterprise autonomy at scale. Together, these three capabilities form an opportunity triad that redefines how autonomous enterprises are built and operated. Combined, they allow AI systems to execute decisions inside real business workflows, operate across the enterprise, and remain viable under cost, governance, and operational constraints.

Opportunity 1: Agentic AI

74% of enterprises plan to deploy AI agents within two years, marking a shift from AI that assists to AI that executes. In autonomous enterprises, agentic AI enables systems to reason, decide, and act continuously inside live business workflows.

Result:
Execute decisions end-to-end with minimal human bottlenecks.

Opportunity 2: Knowledge platform

A knowledge platform unifies structured data, unstructured knowledge, and live operational signals into a governed business context layer. This ensures that AI decisions are informed and executable inside active business processes.

Result:
Execute decisions with real-time business context.

Opportunity 3: Hybrid AI stack

A hybrid AI stack allows AI systems to execute seamlessly across cloud and onprem environments. This allows autonomous AI to run where enterprise data, processes, and regulatory constraints require.

Result:
Execute AI anywhere without sacrificing control, performance, or flexibility.

5 ‘C’s’ of enterprise autonomy

Enterprise autonomy requires AI systems that can execute reliably at scale. The five ‘C’s’ show how the three structural barriers are addressed while simultaneously leveraging the opportunity triad. Together, they provide a clear blueprint for how enterprise AI systems must be designed to operate autonomously and deliver modern AI outcomes.

1. **Combine:** Bring all AI capabilities—data, analytics, ML, models, agents, and applications—into a single, unified system.

2. **Context:** Embed AI with full enterprise business context, spanning structured data, unstructured knowledge, and realtime signals.

3. **Control:** Govern AI execution with built-in guardrails across security, cost, performance, and lifecycle management.

4. **Consume:** Make AI outcomes easy to access and use across the enterprise, enabling any skill level with low-code speed and full-code control.

5. **Connect:** Integrate AI with open ecosystems—models, frameworks, tools, and runtimes, without vendor lockin.

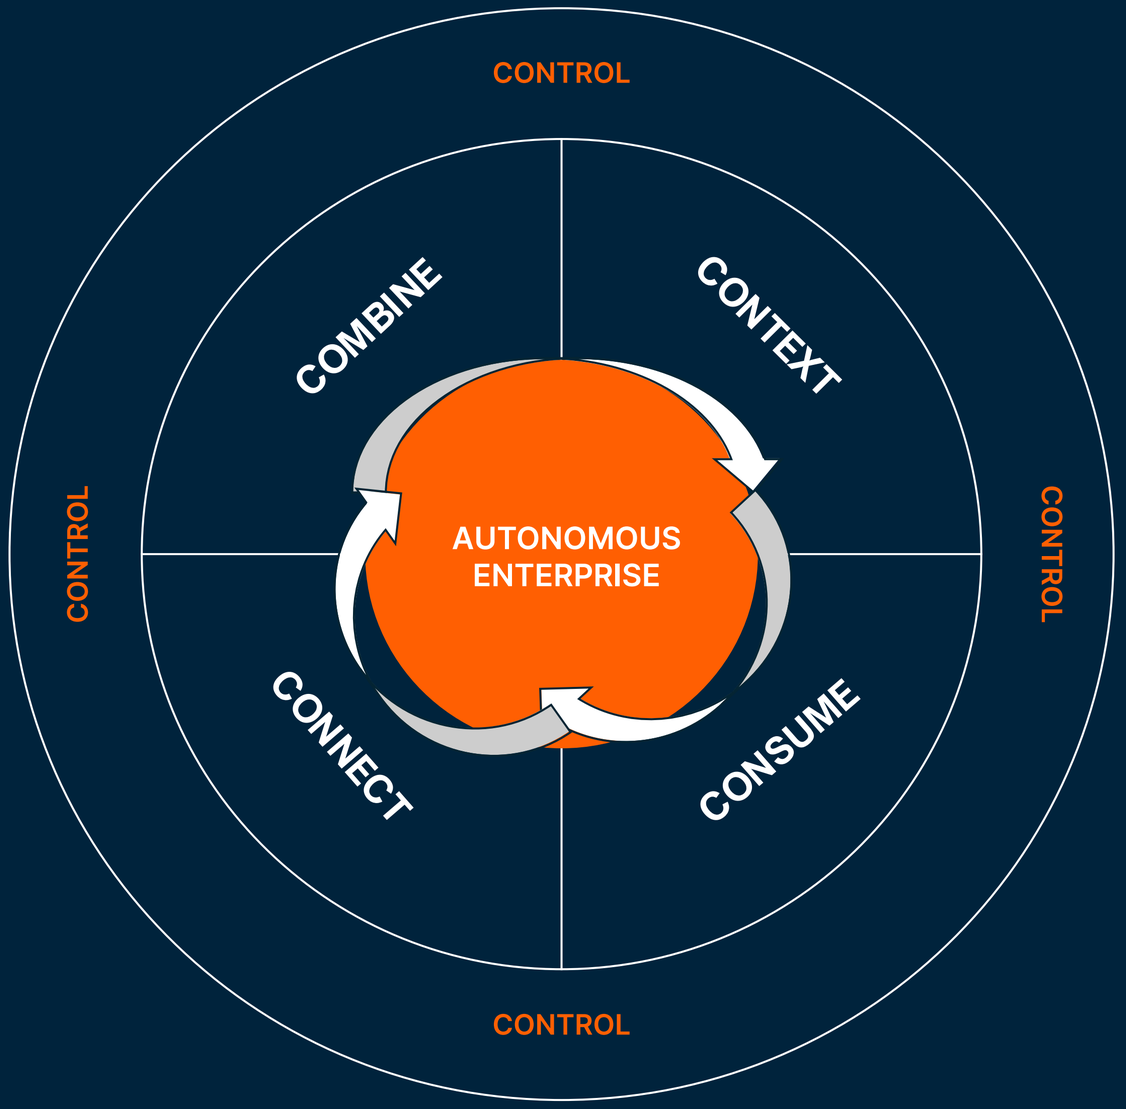

The solution: Teradata AI Studio

AI Studio is a single place where users and creators build, activate, and manage AI outcomes at scale

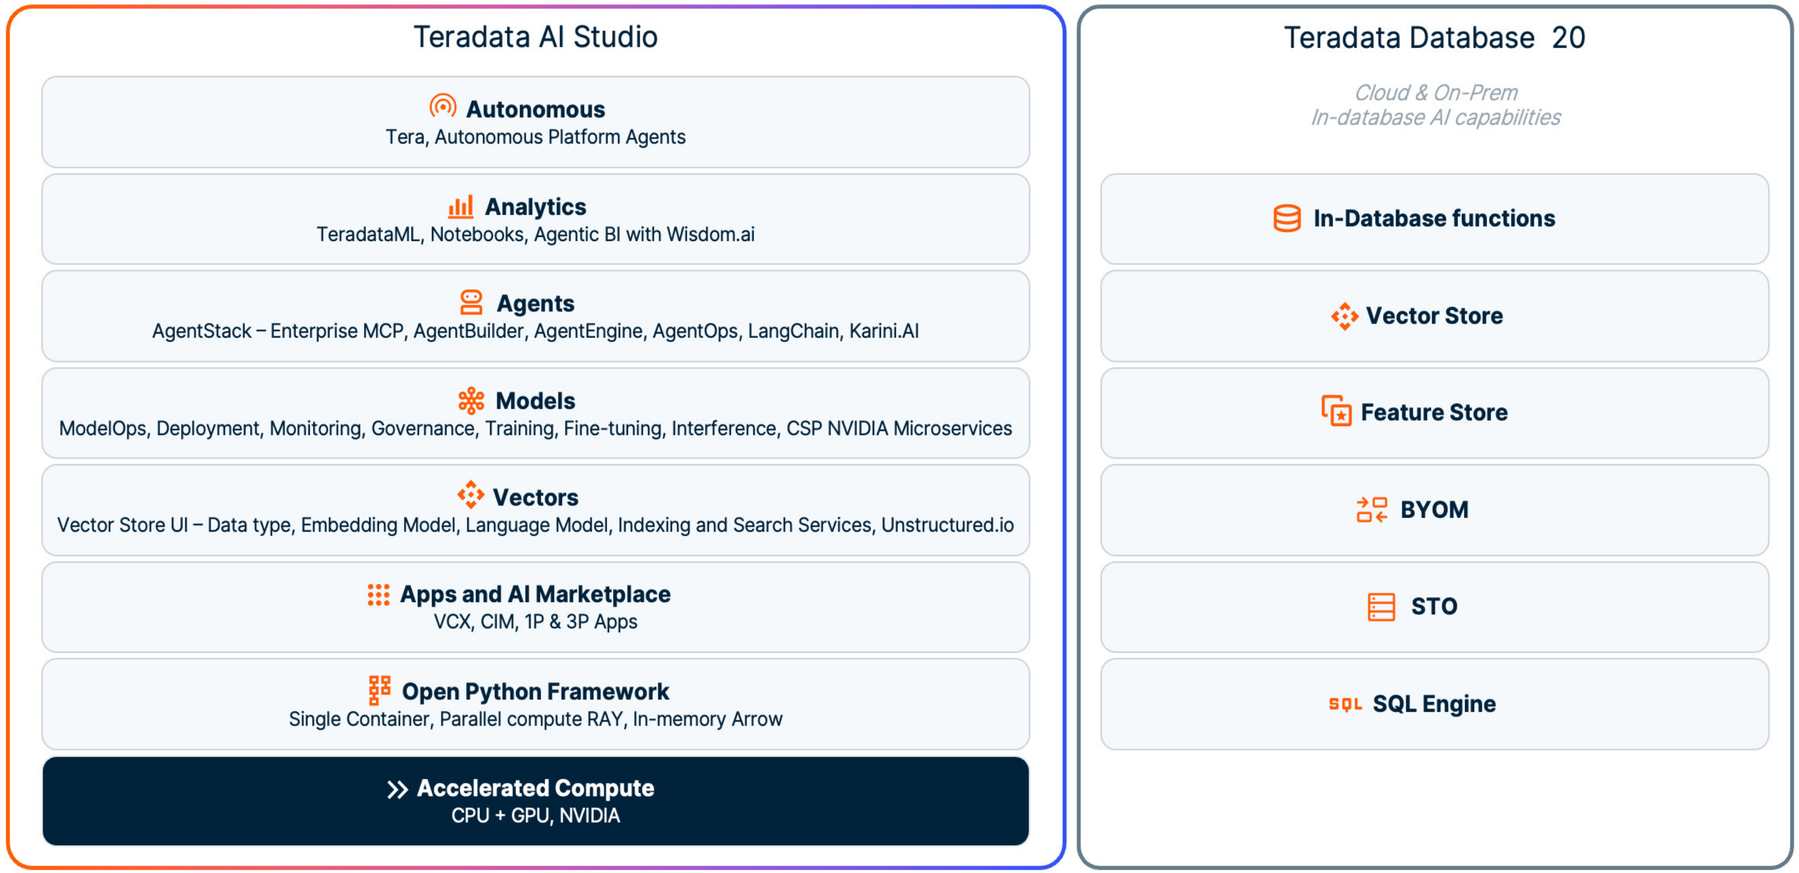

Technology Pillars of AI Studio

Autonomous enterprises are built on AI systems that can execute continuously and reliably at scale. Achieving this requires more than individual AI capabilities. It requires a complete execution foundation that turns modern AI into real, repeatable business outcomes.

The technology pillars of Teradata AI Studio provide this execution foundation. Each pillar solves a specific requirement for autonomous execution by unifying AI systems, embedding enterprise context, enforcing governance, enabling hybrid connectivity, and making AI outcomes usable across the organization.

On their own, these pillars address distinct execution needs. Together, they allow AI to move beyond isolated experiments and operate as a dependable enterprise execution system, delivering modern AI outcomes and advancing the enterprise toward autonomy.

AUTONOMOUS
Capabilities for autonomous AI, including a conversational, task-oriented assistant embedded in the Teradata AI Studio experience.

AGENTS
End to end agent development that enables organizations to move from isolated pilots to production grade agents by simplifying the entire agent lifecycle.

ANALYTICS
Powerful analytics capabilities powered by Teradata ML for end-to-end AI/ML, pre-configured notebooks, and advanced capabilities such as graphic analytics and agentic BI.

MODELS
Comprehensive model lifecycle management that takes AI from experimentation to enterprise scale production.

VECTORS
Vector store UI for unstructured data ingestion covering all steps from embedding to search and retrieval.

APPS AND AI MARKETPLACE
Apps, such as customer experience (CX), transform how organizations engage with customers by delivering intelligent, real-time personalization throughout the entire customer journey.

OPEN PYTHON FRAMEWORK
Offers a flexible, enterprise grade compute purpose built for data scientists and developers to harness the full power of opensource Python libraries such as RAY and Arrow framework.

Why Teradata AI Studio

1

Unified platform experience

Integrate analytics, vector intelligence, model management, and agentic automation under one control plane

Simplifies architecture, accelerates adoption, ensures consistent governance

2

Trusted and governed by design

Extend Teradata's proven governance — identity, lineage, audit, and compliance — across all AI assets

Enables explainable, accountable, policy-aligned AI operations

3

Hybrid and sovereign by default

Deploy seamlessly in cloud, on premises, or AI Factory with identical functionality

Addresses residency, cost predictability, regulatory mandates

4

Open and extensible

Integrate external models and frameworks through open APIs

Provides flexibility and choice while keeping governance centralized

5

Agentic and autonomous intelligence

Enable domain-aware, self-operating agents that execute within enterprise policy boundaries

Moves from predictive analytics to actionable automation

6

Enterprise-scale economics

Deliver superior price-performance through Teradata's MPP architecture, workload management, and shared compute

Supports mission-critical AI workloads with predictable cost

Six pillars, one vision:
To deliver the most flexible, scalable, and autonomous AI, driving breakthrough intelligence and industry best ROI for AI outcomes

Modern AI outcomes in practice

teradata.

Defense intelligence

Business problem

The rules of warfare are being rewritten by cost asymmetry. **A main battle tank—representing $15 million in mobile firepower and crew protection—can now be neutralized by a $300 drone.** This imbalance fundamentally favors cheap, disposable platforms over expensive, irreplaceable assets.

Modern AI outcome

Troops capture images of camouflaged assets and AI instantly assesses concealment effectiveness—delivering real-time guidance to improve survivability.

This closes the loop between detection risk and survivability, transforming a $15 million asset's defense to intelligence-driven protection.

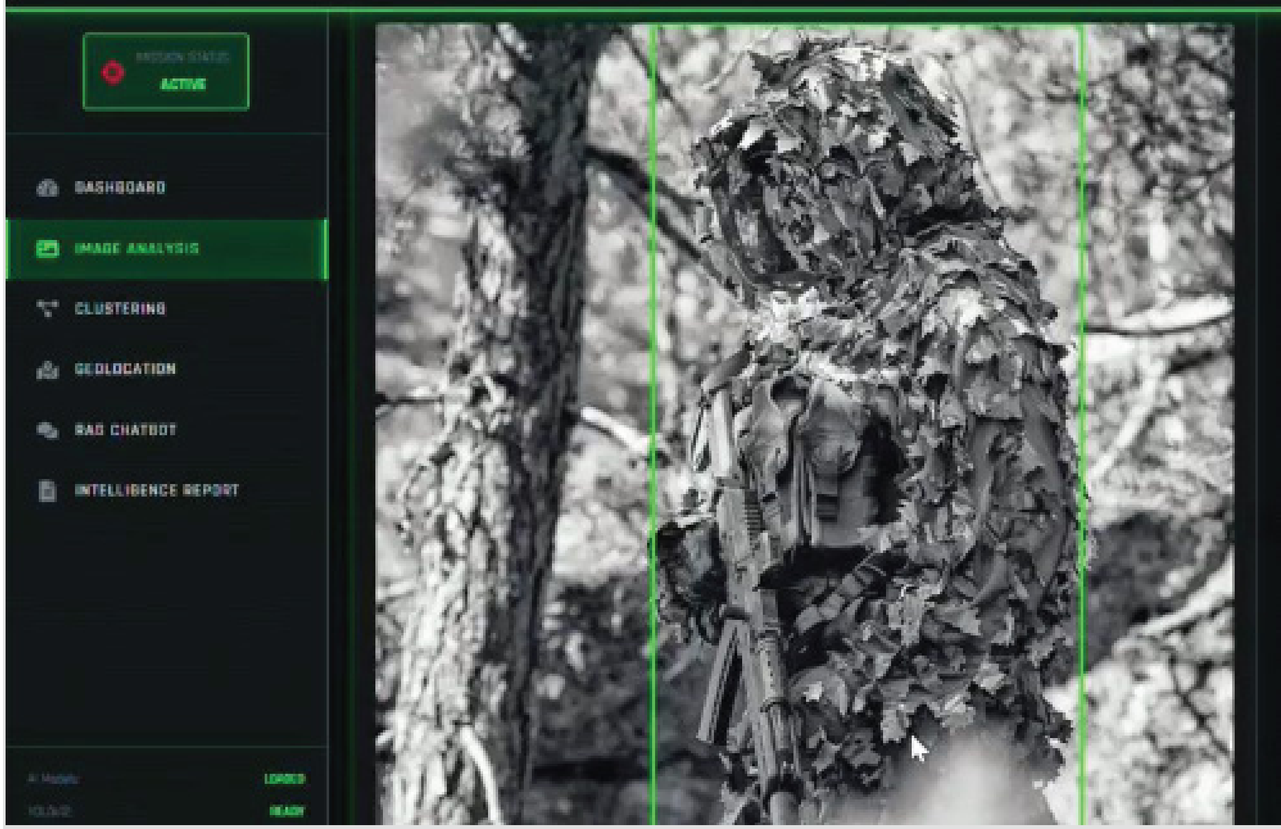

AI Capabilities required

- **Agents** autonomously analyze incoming battlefield imagery, reason over terrain and threat context, and deliver real-time concealment recommendations to soldiers.
- **Models** ensure detection and computer vision models are continuously retrained, versioned, and deployed as battlefield conditions and adversarial tactics evolve.
- **Vectors** enable ingestion, embedding, indexing, and retrieval of multi-modal data including terrain imagery, threat patterns, and concealment best practices.
- **Open PythonFramework** provides parallel compute to train computer vision models and perform real-time image inference using open-source Python libraries

An AI agent for every banking customer

Business problem

The rules of financial security are being shaped by an advice gap. Fewer than one in 10 adults ever receive independent professional financial advice, leaving the majority to make critical pension and savings decisions alone. **Without scalable access to financial guidance, developed economies face growing pressure on public systems as private retirement outcomes continue to fall short.**

Modern AI outcome

The solution delivers a **personal financial-advice agent to every banking customer.** These AI agents continuously monitor customer behavior in real time, helping individuals to make better day-to-day financial decisions. Expert investment agents sit behind them, providing personalized, life-stage-appropriate investment recommendations. **This scales professional-grade financial advice—previously available to only the wealthiest 10%—to millions of customers**

AI Capabilities required

- **Agents** continuously monitor customer behavior and life-stage signals to deliver personalized, real-time financial guidance to each banking customer
- **Models** ensure customer-level prediction and recommendation models are continuously updated to reflect changing behavior, market conditions, and risk profiles
- **Vectors** enable next-best-action intelligence by analyzing transactional, behavioral, and financial signals to support personalized advice
- **Open PythonFramework** provides scalable, parallel compute to train and operate customer-level models, enabling personalized financial intelligence for millions

Logistic optimization using AI agents

Business problem

Most of your purchases—including your phone, your clothes, and the food on your table—probably traveled on a container ship at some point.

But the industry that moves all of this is still shockingly manual and creates problems that can cost billions.

Modern AI outcome

The solution delivers an **agentic approach** to the vessel re-routing problem. AI agents leverage **machine learning** and **knowledge models** to identify optimal routes that avoid port congestion.

Before executing any re-routing, the solution evaluates the potential impact on cost, schedule, and fuel consumption. These insights enable operators to make informed decisions before taking the final re-routing action.

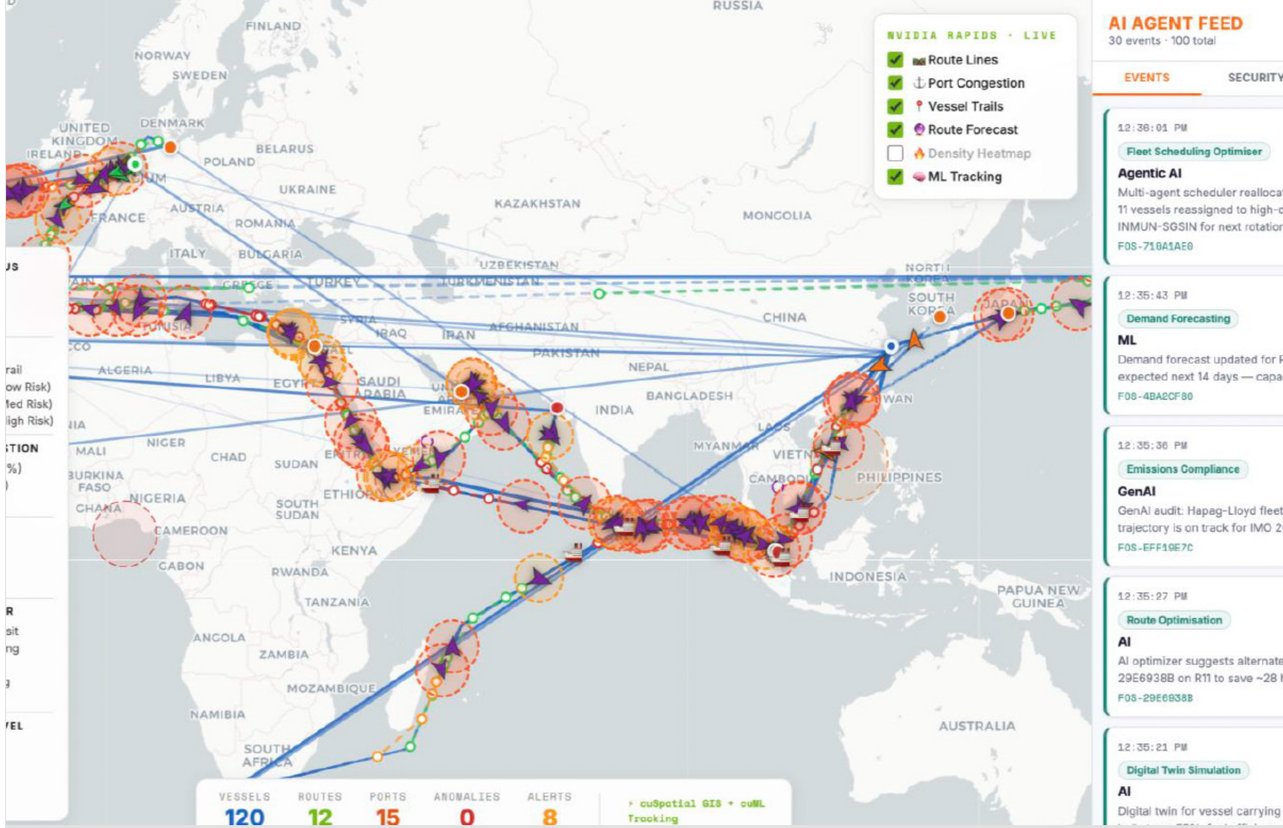

AI Capabilities required

- **Agents** continuously monitor vessel routes, operational constraints, and network conditions to generate real-time re-routing recommendations
- **Models** ensure ETA prediction and route optimization models are continuously updated to reflect changing traffic patterns, port conditions, and execution outcomes
- **Vectors** analyze historical and real-time logistics signals like port congestion, dwell times, and throughput and support optimal routing decisions
- **Open PythonFramework** provides scalable, parallel compute to run machine learning, geospatial, and network graph analytics required to optimize logistics routes at scale

teradata.